OPEN
Compute Project

Datacenter NVMe® SSD Specification

Version 2.5 (09282023)

Author: Ross Stenfort, Meta
Author: Lee Prewitt, Microsoft
Author: Paul Kaler, HPE
Author: David Black, Austin Bolen, Dell Technologies
Author: Chris Sabol, Charles Kunzman, Google

Table of Contents

1 OCP

1.1 License OWF Option

Contributions to this Specification are made under the terms and conditions set forth in Open Web Foundation Contributor License Agreement (“OWF CLA 1.0”) (“Contribution License”) by:

Meta
Microsoft
HPE
Dell
Google

Usage of this Specification is governed by the terms and conditions set forth in the Open Web Foundation Final Specification Agreement (“OWFa 1.0”).

Note: The following clarifications, which distinguish technology licensed in the Contribution License and/or Specification License from those technologies merely referenced (but not licensed), were accepted by the Incubation Committee of the OCP:

NONE

1.2 System Firmware

All products seeking OCP Accepted™ Product Recognition must complete the Open System Firmware (OSF) Tab in the 2021 Supplier Requirements Checklist. The completed checklist shall be uploaded and available at: *https://github.com/opencomputeproject/OpenSystemFirmware/[vendor_name]/[product_name]/*

Note to author: replace [vendor_name] and [product_name] with actual company name and product identifier.

1.3 Hardware Management

1.3.1 Compliance

All products seeking OCP Inspired™ or OCP Accepted ™ Product Recognition shall comply with the *OCP Hardware Management Baseline Profile V1.0* and provide such evidence by completing the Hardware Management Tab in the 2021 Supplier Requirements Checklist.

1.3.2 BMC Source Availability (if applicable)

All Products seeking OCP Accepted™ Product Recognition shall have source code and binary blobs submitted for BMC, if applicable. The BMC management source code shall be uploaded at: *https://github.com/opencomputeproject/Hardware Management/[vendor_name]/[product_name]/*.

1.4 Security

All products seeking OCP Inspired™ or OCP Accepted™ Product Recognition shall have a completed Security Profile in the 2021 Supplier Requirements Checklist. Whether the answer is a yes or no, the profile must be completed. For Additional Security Badges (Bronze/Silver/Gold), please fill out the Security Profile in accordance with the requirements for that level. Security Badges will be reassessed on an annual basis as requirements are subject to change.

2 Overview

This document is to define the requirements for a Datacenter NVMe SSD (DSSD) for use in datacenters.

3 Scope

This document covers PCIe-attached SSDs using NVM Express.

4 NVM Express Requirements

4.1 Overview

Requirement ID	Description
NVMe-1	The device shall comply with all required features of the NVM Express Base 2.0 Specification, the NVM Command Set 1.0 Specification and the NVMe over PCIe Transport 1.0 Specification, or later except as overridden by this specification. Optional features in those three specifications shall be implemented per the requirements of this specification. The vendor shall provide an NVMe compliance report that demonstrates compliance.
NVMe-2	Any optional NVMe features supported by the device not described in this document shall be clearly documented and disclosed to the customer.
NVMe-3	Any vendor unique features supported by the device not described in this document shall be clearly documented and disclosed to the customer.

4.2 NVMe Reset Supported

Requirement ID	Description
NVMeR-1	NVM Subsystem Reset (NSSR) shall be supported via the NSSRC register.
NVMeR-2	NVMe Controller Reset (CC.EN cleared to 0b) shall be supported.

4.3 NVMe Controller Configuration and Behavior

Requirement ID	Description
NVMe-CFG-1	The default arbitration shall be Round-Robin. Weighted Round Robin with Urgent Class Priority shall be supported.
NVMe-CFG-2	The device shall support a Maximum Data Transfer Size (MDTS) value of at least 256KB.
NVMe-CFG-3	The device firmware shall support the reporting of CSTS.CFS.
NVMe-CFG-4	Obsolete. Replaced by LABL-11.
NVMe-CFG-5	The minimum supported queue depth shall be 1024 per submission queue.
NVMe-CFG-6	The minimum number of I/O Queue Pairs shall be 512.
NVMe-CFG-7	The single default namespace (see SECTOR-3) shall have a unique non-zero value in the EUI64 (Extended Unique Identifier) field in the Identify Namespace data structure.

Requirement ID	Description
	For each namespace that is created other than the initial namespace configured in the factory, the device shall clear the EUI64 field in the Identify Namespace data structure to 0h.
NVMe-CFG-8	The device shall support a non-zero NGUID per Namespace that is never reused (i.e., the UIDREUSE bit in the Common Namespace Features field of the Identify Namespace data structure shall be set to ‘1’).
NVMe-CFG-9	The Doorbell Stride field in the Controller Capabilities data structure (CAP.DSTRD) shall be cleared to zero.
NVMe-CFG-10	The device shall place all Doorbells in a single contiguous 4096-byte area of MMIO space that is aligned to a 4096-byte boundary.
NVMe-CFG-11	If the Controller Memory Buffer (CMB) is supported, the implementation shall comply with the NVM Express Base 2.0 Specification or later (see NVMe-1). This includes the requirement that the CMB shall be disabled by default. CMB implementations that comply with any version of the NVM Express Base Specification prior to version 1.4 are prohibited.
NVMe-CFG-12	The Controller Power Scope field in the Controller Capabilities property shall be set to a value of 11b to indicate it is NVM subsystem scope (i.e., the NVM subsystem does not support multiple domains).
NVMe-CFG-13	The device shall support the NVM Command Set (i.e., the Command Sets Supported field in the Controller Capabilities property shall set bit 0 to 1b and the Controller Type field in the Identify Controller data structure shall be set to 01h).
NVMe-CFG-14	The device shall support a maximum Submission Queue Entry Size of 64 bytes (i.e., bits 7:4 of the SQES field in the Identify Controller data structure shall be set to 6).
NVMe-CFG-15	The device shall support a maximum Completion Queue Entry Size of 16 bytes (i.e., bits 7:4 of the CQES field in the Identify Controller data structure shall be set to 4).
NVMe-CFG-16	The device shall clear the Memory Page Size Minimum (MPSMIN) field to 0h (4 KiB minimum page size).
NVMe-CFG-17	The device shall set the Memory Page Size Maximum (MPSMAX) field to a value of 1h or greater (8 KiB or greater maximum page size).
NVMe-CFG-18	The device shall support 256 Error Log Page Entries.
NVMe-CFG-19	The device shall clear the Host Memory Buffer Minimum Size field and the Host Memory Buffer Preferred Size field in the Identify Controller data structure to 0h (i.e., the Host Memory Buffer feature shall not be supported).
NVMe-CFG-20	The device shall clear the CAP.CRMS.CRIMS bit to 0b.

4.4 NVMe Admin Command Set

The device shall support the following mandatory and optional NVMe Admin commands:

Requirement ID	Description
NVMe-AD-1	The device shall support all mandatory NVMe Admin commands.

Requirement ID	Description
NVMe-AD-2	Identify – In addition to supporting all the mandatory CNS values and the associated mandatory fields within the CNS, the following optional fields in the CNS shall be supported: • Format progress indicator (FPI) with an update granularity of 1%. • I/O Performance and Endurance Hints: o NSFEAT bit 4 = 1b.
NVMe-AD-3	Namespace Management command shall be supported.
NVMe-AD-4	Namespace Attachment command shall be supported.
NVMe-AD-5	Format NVM command shall be supported. Secure Erase Settings (SES) values 000b, 001b and 010b shall be supported.
NVMe-AD-6	The device shall support the NVMe-MI Receive and NVMe-MI Send commands (i.e., bit 6 in the Optional Admin Command Support field in the Identify Controller data structure shall be set to 1b). The device shall support all NVMe-MI commands that map to the NVMe-MI Receive command and are mandatory using the in-band tunneling mechanism. The device shall prohibit all NVMe-MI commands that map to the NVMe-MI Send command.
NVMe-AD-7	The device shall support the Sanitize command. Block Erase (010b), Overwrite (011b) and Crypto Erase (100b) sanitize operations shall be supported.
NVMe-AD-8	Once data is purged (see SEC-43) due to: • a Sanitize command that has the No-Deallocate After Sanitize bit cleared to 0b; or • a Format NVM command that specifies User Data Erase or Cryptographic Erase; or • a TCG Revert method, RevertSP method or GenKey method. The device shall enable reads to the purged LBAs without generating errors. Once a purge completes, the device shall return all zeroes for host reads to LBAs that have not been written since the last purge completed.
NVMe-AD-9	Obsolete. Incorporated into NVMe-AD-8.
NVMe-AD-10	The device shall support Identify command UUID List functionality (CNS value 17h).
NVMe-AD-11	If the Namespace Identifier (NSID) is not used for a specific NVMe Admin command, and the host specifies a non-zero NSID, then the controller shall abort the command with status Invalid Field in Command.
NVMe-AD-12	The Device Self-test command shall be supported.
NVMe-AD-13	Obsolete. See SELFTST-8.
NVMe-AD-14	The device shall support the Lockdown command. For device profile specific requirements see DP-CFG-11.
NVMe-AD-15	The Lockdown command shall support locking down all supported Admin commands on the in-band interface including the Lockdown command itself and all Admin commands not referenced in this specification.
NVMe-AD-16	The Lockdown command shall support locking down all supported Features on the in-band interface, including those not referenced in this specification.
NVMe-AD-17	The device shall set the No-Deallocate Modifies Media After Sanitize field in the Sanitize Capabilities in the Identify Controller data structure to 01b. Additional media modification after Sanitize is prohibited.

Requirement ID	Description
NVMe-AD-18	The device shall clear the No-Deallocate Inhibited bit in the Sanitize Capabilities in the Identify Controller data structure to 0b.
NVMe-AD-19	If the device cannot return zeroes after a sanitize operation successfully completes that had been started by a Sanitize command with the No-Deallocate After Sanitize bit set to 1b, then the device shall instead perform a deallocation even though it was requested not to deallocate and shall set the Sanitize Status field to 100b in the Sanitize Status Log Page.
NVMe-AD-20	Obsolete. Incorporated into NVMe-AD-8.
NVMe-AD-21	The device shall not write data to the media as part of a Crypto Erase sanitize operation, Block Erase sanitize operation, Format NVM Cryptographic Erase, Format NVM User Data Erase, or TCG Opal Revert/RevertSP/GenKey method.
NVMe-AD-22	The device shall be shipped from the factory with the Sanitize Status field in the Sanitize Status log page cleared to 000b indicating the NVM subsystem has never been sanitized and the Global Data Erased bit in the Sanitize Status field set to 1b.
NVMe-AD-23	The device shall support the Format NVM command on a per namespace basis (i.e., bit 0 and bit 1 of the Format NVM Attributes field in Identify Controller shall be cleared to 00b).
NVMe-AD-24	The device shall report a minimum of 4 in the Asynchronous Event Request Limit field in the Identify Controller data structure.
NVMe-AD-25	If asynchronous events occur for which reporting is enabled and there are no Asynchronous Event Request commands outstanding, the controller shall retain the event information for those Asynchronous Event Types and use that information for responses to future Asynchronous Event Request commands. Queued duplicate Asynchronous events shall be reported once.

4.4.1 Namespace Management/Attachment Commands

The namespace management command along with the attach/detach commands is used to increase device over-provisioning beyond the default minimum over-provisioning.

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>NSM-1</td><td>The namespace management commands shall be supported on all namespaces.</td></tr>
<tr><td>NSM-2</td><td>When creating a namespace, the default “Formatted LBA Size” parameter (FLBAS = 0) in the Identify Namespace Data Structure (byte 26) shall correspond to the default sector size set at the factory.</td></tr>
<tr><td>NSM-3</td><td>When formatting the device with the Format NVM command, the default “LBA Format” parameter (LBAF = 0) in Command Dword 10 bits 3:0 shall correspond to the default sector size set at the factory.</td></tr>
<tr><td>NSM-4</td><td>The device shall support a minimum of 16 Namespaces (see Section 13 Device Profiles).</td></tr>
<tr><td>NSM-5</td><td>For some models (see Section 13 Device Profiles), the minimum number of namespaces shall be at least 128 Namespaces at 7TB based on 16 Namespaces per TB of usable capacity. The number of Namespaces shall be based on the device usable capacity as follows:
<table>
<tr><th>Device Usable Capacity</th><th>Minimum Number of Namespaces</th></tr>
<tr><td><= 1TB</td><td>16</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td><table>
<tr><td>> 1TB but <= 2TB</td><td>32</td></tr>
<tr><td>> 2TB but <= 3TB</td><td>48</td></tr>
<tr><td>> 3TB but <= 4TB</td><td>64</td></tr>
<tr><td>…</td><td>…</td></tr>
<tr><td>> 7TB</td><td>128</td></tr>
</table></td></tr>
<tr><td>NSM-6</td><td>The device shall report at least one Namespace Granularity Descriptor in the Namespace Granularity List.</td></tr>
<tr><td>NSM-7</td><td>The device shall support the TNVMCAP and UNVMCAP fields in the Identify Controller Data Structure.</td></tr>
</table>

4.4.2 Namespace Utilization (NUSE)

Requirement ID	Description
NUSE-1	The NUSE shall be equal to the number of logical blocks currently allocated in the namespace. NUSE cannot be hardcoded to be equal to NCAP. Here is an example for a 200GB device: 1. After a Format NVM command User Data Erase (SES = 001b), NUSE would be zero. And the usage data would reflect that: 0.00GB. 2. After writing 1 GB worth of data, the usage data is expected to show the following: 1.00GB. 3. After filling the device, the usage data is expected to show the following: 200.00GB. 4. If the host issues a 10GB de-allocate command and the device completes de-allocation of the data, the usage data would show the following: 190.00GB.

4.4.3 UUID for DSSD Specific Information

A UUID has been defined for use in commands to ensure that the vendor specific Log Identifiers and Feature Identifiers used in this specification access the functionality defined in this specification (i.e., do not access other vendor specific functionality that may use the same vendor specific identifiers).

Requirement ID	Description
UUID-1	The UUID List (NVMe-AD-10) shall contain a UUID List Entry that contains the UUID value C194D55BE0944794A21D29998F56BE6Fh. The Identifier Association field in that UUID List Entry shall be cleared to 00b. For clarity: UUIDDw0 (bytes 51:48) = 0xC194D55B UUIDDw1 (bytes 55:52) = 0xE0944794 UUIDDw2 (bytes 59:56) = 0xA21D299 UUIDDw3 (bytes 63:60) = 0x8F56BE6F
UUID-2	The Get Features and Set Features commands shall support UUID Index functionality.
UUID-3	A Get Features command or a Set Features command with: • the UUID Index of the UUID (UUID-1) in the UUID List (NVMe-AD-10) or a zero UUID Index; and

Requirement ID	Description
	• a vendor-specific Feature Identifier that is used in this specification (see Section 4.15 Set/Get Feature Requirements) shall access the vendor specific Feature defined in this specification.
UUID-4	The Get Log Page command shall support UUID Index functionality.
UUID-5	A Get Log Page command with: • the UUID Index of the UUID (UUID-1) in the UUID List (NVMe-AD-10) or a zero UUID Index; and • a vendor-specific Log Page Identifier that is used in this specification (see Section 4.8 Log Page Requirements) shall access the vendor specific Log Page defined in this specification.
UUID-6	The vendor-specific events in the Persistent Event log page shall support UUID Index functionality.

4.5 NVMe I/O Command Set

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>NVMe-IO-1</td><td>The device shall support all mandatory NVMe I/O commands.</td></tr>
<tr><td>NVMe-IO-2</td><td>The device shall support the Dataset Management command. For additional information see TRIM-1.</td></tr>
<tr><td>NVMe-IO-3</td><td>Since the device is power fail safe (e.g., has Power Loss Protection (PLP)) the performance shall not be degraded by any of the following:
• FUA – i.e., forced unit access shall not incur a performance penalty.
• Flush Cache – i.e., flush cache shall have no effect as the PLP makes any cache non-volatile.
• Volatile Write Cache (Feature Identifier 06h) Set Feature to disable write-cache. This command shall fail as described in the NVMe Standard 1.4b as there is no volatile write cache.</td></tr>
<tr><td>NVMe-IO-4</td><td>The device shall support the Write Zeroes command. The following bits of the Write Zeroes command shall be supported:
• De-allocate (DEAC) bit.
• Force Unit Access (FUA) bit.</td></tr>
<tr><td>NVMe-IO-5</td><td>The Write Zeroes command shall have the following behavior:
<table>
<tr><th>DEAC</th><th>FUA</th><th>Behavior</th></tr>
<tr><td>0b</td><td>0b</td><td>The device shall follow the NVMe NVM Command Set Specification, see NVMe-1.</td></tr>
<tr><td>0b</td><td>1b</td><td>The device shall follow the NVMe NVM Command Set Specification, see NVMe-1.</td></tr>
<tr><td>1b</td><td>1b</td><td>The device shall follow the NVMe NVM Command Set Specification, see NVMe-1.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td><table><tr><td>1b</td><td>0b</td><td>See NVMe-IO-6.</td></tr></table></td></tr>
<tr><td>NVMe-IO-6</td><td>If the Write Zeroes DEAC bit is set to 1b and the FUA bit is cleared to 0b, the device shall deallocate the specified blocks and shall return a zero value for any subsequent read to the specified blocks until modified by a command (e.g., Write, Copy, etc.), regardless of the behavior of the Dataset Management command.</td></tr>
<tr><td>NVMe-IO-7</td><td>With the DEAC bit set to 1b and the FUA bit cleared to 0b one or more Write Zeroes command(s) shall be able to completely deallocate the entire device including updating the FTL map in less than one minute.</td></tr>
<tr><td>NVMe-IO-8</td><td>The device shall support the Compare command.</td></tr>
<tr><td>NVMe-IO-9</td><td>The device shall support the Compare and Write fused command pair.</td></tr>
<tr><td>NVMe-IO-10</td><td>For some models (see Section 13 Device Profiles), the device shall support the Write Uncorrectable command.</td></tr>
<tr><td>NVMe-IO-11</td><td>The Write Uncorrectable command shall support marking LBAs uncorrectable at a single LBA granularity regardless of the number of LBAs in the FTL indirection granularity.</td></tr>
<tr><td>NVMe-IO-12</td><td>The device shall not limit the number of LBAs that the host is able to specify in a Write Uncorrectable command beyond the minimum and maximum allowed by NVMe. The host shall be able to send a single LBA.</td></tr>
<tr><td>NVMe-IO-13</td><td>There shall be no limit on the total media capacity that can be marked uncorrectable by Write Uncorrectable commands.</td></tr>
<tr><td>NVMe-IO-14</td><td>Uncorrectable errors (e.g., read errors) that are a consequence of a prior Write Uncorrectable command shall not be counted in the Smart / Health Information (Log Identifier 02h) Media and Data Integrity Errors field.</td></tr>
</table>

4.6 Optional NVMe Feature Support

The device shall also support the following NVMe features:

Requirement ID	Description
NVMe-OPT-1	Obsolete. Duplicate of STD-LOG-7 and STD-LOG-8.
NVMe-OPT-2	Timestamp (Feature Identifier 0Eh) shall be supported to align the devices internal logs.
NVMe-OPT-3	Obsolete. See TEL-5.
NVMe-OPT-4	The device shall only clear the Timestamp Origin field to 000b in the Timestamp (Feature Identifier 0Eh) on a main power cycle or NVM Subsystem Reset (e.g., NSSR). The device shall not clear the Timestamp Origin field on a power cycle of only AUX power.
NVMe-OPT-5	The device shall never set the Synch field bit to 1b in the Timestamp (Feature Identifier 0Eh).
NVMe-OPT-6	The device shall only generate a Telemetry Log Changed event on an error condition and shall not generate a Telemetry Log Changed event for periodic logging.
NVMe-OPT-7	The device shall report its Indirection Unit (IU) size in the NPWG field in the Identify Namespace data structure.
NVMe-OPT-8	The device shall support sending Telemetry Log Notices (i.e., bit 3 of the Log Page Attributes field shall be set to '1').

Requirement ID	Description
NVMe-OPT-9	The device shall support the Namespace Attribute Notices event and the associated Changed Namespace List log page (i.e., bit 8 of the Optional Asynchronous Events Supported field in the Identify Controller data structure shall be set to 1b).
NVMe-OPT-10	The device shall support the Firmware Activation Notices event (i.e., bit 9 of the Optional Asynchronous Events Supported field in the Identify Controller data structure shall be set to 1b).
NVMe-OPT-11	The device shall support an Asynchronous Event Request Limit of at least 4.
NVMe-OPT-12	The device shall be composed of a single NVM subsystem.

4.7 Command Timeout

Requirement ID	Description
CTO-1	NVMe Admin commands and TCG commands shall take no more than 10 seconds from submission to completion with no additional commands outstanding (QD1). CTO-1 does not apply to the time taken by background operations initiated by the Device Self-test and Sanitize commands. CTO-1 does not apply to Asynchronous Event Request commands.
CTO-2	The only exception to CTO-1 shall be the Format NVM Command for non-Cryptographic Erase operations. Format NVM Cryptographic Erase and the TCG methods Revert, RevertSP and GenKey shall comply with CTO-1.
CTO-3	Once TTR-2 is satisfied, an individual I/O command (QD1) shall take no more than 8 seconds from submission to completion. The device shall not have more than 7 I/Os take more than 2 seconds in one hour.
CTO-4	I/O command processing time shall not be a function of device capacity.
CTO-5	Device supplier shall disclose any I/O scenario that could violate the timeout requirements in CTO-1 through CTO-4.
CTO-6	The device shall set MDTS, the attributes of the Number of Queues feature (Feature Identifier 07h) and Maximum Queue Entries Supported such that CTO-3 cannot be violated.
CTO-7	A Crypto Erase sanitize operation and any background activity initiated by a TCG command (e.g., as indicated by OutstandingData) shall complete within 10 seconds.
CTO-8	A Block Erase sanitize operation and Format NVM command performing a User Data Erase shall complete within 10 seconds per TB of device capacity (e.g., Block Erase sanitize of a 4TB device shall complete within 40 seconds).
CTO-9	The device shall report the greater of the Sanitize Crypto Erase, Format NVM Cryptographic Erase, TCG Opal Revert, TCG Opal RevertSP, or TCG Opal GenKey time in the Estimated Time for Crypto Erase field in the Sanitize Status log page.
CTO-10	The device shall not set the Estimated Time fields in the Sanitize Status log page to a value of FFFFFFFFh for supported Sanitize operations.
CTO-11	Formats between LBAF values with the Secure Erase Settings field cleared to 000b (No secure erase operation requested) or 2h (Cryptographic Erase) shall complete within 30 seconds.

4.8 Log Page Requirements

4.8.1 Standard Log Page Requirements

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>STD-LOG-1</td><td>The device shall support the Error Information (Log Identifier 01h) log page.</td></tr>
<tr><td>STD-LOG-2</td><td>The device shall support the SMART / Health Information (Log Identifier 02h) log page.</td></tr>
<tr><td>STD-LOG-3</td><td>Under no conditions shall the Percentage Used field in the SMART / Health Information (Log Identifier 02h) log page decrease to a lower value than previously reported.</td></tr>
<tr><td>STD-LOG-4</td><td>The Percentage Used field in the SMART / Health Information (Log Identifier 02h) log page shall be based on the average P/E cycle of the device. In addition, this field shall be based on the actual P/E cycle count of the media and not on the Power on Hours (POH) of the device.</td></tr>
<tr><td>STD-LOG-5</td><td>The device shall support the Firmware Slot Information (Log Identifier 03h) log page.</td></tr>
<tr><td>STD-LOG-6</td><td>The device shall support the Commands Supported and Effects (Log Identifier 05h) log page.</td></tr>
<tr><td>STD-LOG-7</td><td>The device shall support the Telemetry Host-Initiated (Log Identifier 07h) log page.</td></tr>
<tr><td>STD-LOG-8</td><td>The device shall support the Telemetry Controller-Initiated (Log Identifier 08h) log page.</td></tr>
<tr><td>STD-LOG-9</td><td>The device shall support the Persistent Event Log (Log Identifier 0Dh) log page.</td></tr>
<tr><td>STD-LOG-10</td><td>The device shall support the following Persistent Event Log types:
<table>
<tr><th>Type</th><th>Event</th></tr>
<tr><td>01h</td><td>SMART / Health Log Snapshot</td></tr>
<tr><td>02h</td><td>Firmware Commit</td></tr>
<tr><td>03h</td><td>Timestamp Change</td></tr>
<tr><td>04h</td><td>Power-on or Reset</td></tr>
<tr><td>05h</td><td>NVM Subsystem Hardware Error
The device shall not log PCIe Unsupported Request errors that are Advisory Non-Fatal</td></tr>
<tr><td>06h</td><td>Change Namespace</td></tr>
<tr><td>07h</td><td>Format NVM Start</td></tr>
<tr><td>08h</td><td>Format NVM Completion</td></tr>
<tr><td>09h</td><td>Sanitize Start</td></tr>
<tr><td>0Ah</td><td>Sanitize Completion</td></tr>
<tr><td>0Ch</td><td>Telemetry Log Created</td></tr>
<tr><td>0Dh</td><td>Thermal Excursion</td></tr>
<tr><td>DEh</td><td>Vendor Specific
This specification defines vendor specific events to report TCG command execution (see Section 4.8.2 and Section 4.8.3).</td></tr>
</table>
</td></tr>
<tr><td>STD-LOG-11</td><td>The device shall support the LBA Status Information (Log Identifier 0Eh) log page.</td></tr>
<tr><td>STD-LOG-12</td><td>The following behavior types are used to define the functional boundary condition behaviors for saturating or wrapping, under reset conditions and power cycles.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Behavior Type</th><th>Saturating Counter[1]</th><th>Reset[2] Persistent</th><th>Power Cycle/PERST# Persistent</th></tr>
<tr><td>0</td><td colspan="3">Reserved</td></tr>
<tr><td>1</td><td>No (Runtime Value)</td><td>No (Runtime Value)</td><td>No (Runtime Value)</td></tr>
<tr><td>2</td><td>No</td><td>Yes</td><td>Yes</td></tr>
<tr><td>3</td><td>Yes</td><td>Yes</td><td>No</td></tr>
<tr><td>4</td><td>Yes</td><td>Yes</td><td>Yes</td></tr>
<tr><td>5</td><td>Yes</td><td>No</td><td>No</td></tr>
<tr><td>6</td><td>No</td><td>Yes</td><td>No</td></tr>
<tr><td>15-7</td><td colspan="3">Reserved</td></tr>
</table>
[1] See SLOG-2 saturating counters.
[2] All Resets (e.g., Controller, NSSR, FLR, PCIe hot reset, etc.).
For Behavior Types 1 the counter shall reset to zero on either a Reset or a Power Cycle/PERST#. For Behavior Type 6 the counter shall reset to zero only on a Power Cycle/PERST#.
</td></tr>
<tr><td>STD-LOG-13</td><td>The device shall support the Device Self-test (Log Identifier 06h) log page.</td></tr>
<tr><td>STD-LOG-14</td><td>The device shall support the Command and Feature Lockdown (14h) log page.</td></tr>
<tr><td>STD-LOG-15</td><td>The device shall support the Physical Interface Receiver Eye Opening Measurement (Log Identifier 19h) log page per TP4119a.</td></tr>
<tr><td>STD-LOG-16</td><td>In addition to any receiver margining data and image plots, the device shall store information required to evaluate the signal integrity of the receiver such as equalization values like continuous time linear equalization (CTLE), variable gain amplifier (VGA), boost, and decision feedback equalization (DFE) settings in the Eye Data field of the Physical Interface Receiver Eye Opening Measurement log page.</td></tr>
<tr><td>STD-LOG-17</td><td>The device shall support the Endurance Group Information (Log Identifier 09h) log page. The device shall use a single Endurance Group for the entire device.</td></tr>
<tr><td>STD-LOG-18</td><td>The device shall not support the SMART / Health Information log page on a per namespace basis (i.e., bit 0 of the Log Page Attributes field in the Identify Controller data structure shall be cleared to ‘0’).</td></tr>
<tr><td>STD-LOG-19</td><td>The Persistent Event log page, Telemetry Host-Initiated log page, Telemetry Controller-Initiated log page, and Error Information log page shall not contain user data, or the result of any computation performed on user data.</td></tr>
<tr><td>STD-LOG-20</td><td>The Persistent Event log page events shall be ordered so that the more recent events are reported earlier in the log than older events (i.e., the most recent event is Persistent Event 0, and the oldest events are removed first when the log is trimmed due to lack of space to store all events).</td></tr>
<tr><td>STD-LOG-21</td><td>The persistent event log reporting context shall be held until released by the host or until at least 5 minutes has elapsed since the last host access of the log. New events that occur while the context is established shall not be dropped.</td></tr>
</table>

Requirement ID	Description
STD-LOG-22	The maximum size of the Persistent Event log page reported in the Persistent Event Log Size field in the Identify Controller data structure shall be 5 MiB or greater.
STD-LOG-23	The device shall not lose debug data relevant to the current operating state due to any operation. Data that does not contain user data or data that could be used to derive user data including the Telemetry Host-Initiated log page, Telemetry Controller-Initiated log page, Persistent Event log page, Error Information log page, Vendor-specific log pages, SMART / Health Information log page (Log Identifier 02h), and SMART / Health Information Extended (Log Identifier C0h) shall be persisted across power loss, all resets, Format NVM command, purge operations (see SEC-43, or device entering any protected mode (e.g., read only or panic mode). Issues with the device's encryption engine shall not prevent access to the data specified by this requirement.

4.8.2 TCG Activity Events for Persistent Event Log

This defines the requirements for recording TCG Activity Events in the Persistent Event Log.

Requirement ID	Description
TCGHST-PE-1	The device shall support recording at least 100 TCG Activity Events in the Persistent Event Log without deletion of any TCG Activity Events to make room for new events.
TCGHST-PE-2	A TCG Activity Event entry shall be recorded in the Persistent Event Log whenever any of the following TCG activities occurs (the event is recorded upon completion of the activity): • Level 0 Discovery • Start Session • Authenticate method (entries only recorded when the authority specified is not locked out) • Close Session by the host when it sends an End-Of-Session token • Close Session due to timeout • Properties method • Stack Reset • TPer Reset • Get method • Set method • GenKey method • Random method • Activate method • Revert method • RevertSP method • Reactivate method • Erase method • Block SID Authentication • Assign method • Deassign method

Requirement ID	Description
TCGHST-PE-3	Redundant TCG activities shall not generate new TCG Command Event entries in the Persistent Event Log. A TCG activity shall be considered redundant if it meets ALL the criteria below with respect to a TCG Activity Event that is already recorded in the Persistent Event Log: • Power cycle count is the same; and • Active firmware revision is the same; and • TCG activity is the same as the most recent TCG Activity Event entry; and • Tper state remains the same (i.e., there has been no change either in the device behavior or host visible state); and • The Result field is the same. The following are examples of TCG activities that do not generate a new TCG Activity Event entry: • Setting an attribute of a TCG Object to a value that is the same as the current value. • Getting or setting the value of an informative attribute of a TCG object (e.g., the CommonName of a Locking Object). • Getting data from or setting data to byte tables. • An unsuccessful authentication attempt that returns AUTHORITY_LOCKED_OUT status.
TCGHST-PE-4	TCG Activity Event entries shall not be removed for any TCG-specific reason (e.g., TCG Revert/RevertSP or TCG Manufactured-Inactive), except as required to make room in the Persistent Event Log for additional TCG Activity Event entries (see TCGHST-PE-1).
TCGHST-PE-5	TCG Activity Event format shall follow the requirements in the TCG Activity Event Format section below.

4.8.3 TCG Activity Event Format

This defines the TCG Activity Event format for the Persistent Error Log.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCGCE-1	0	Event Type	1	Shall be set to Deh to indicate a vendor-specific event.
TCGCE-2	1	Event Type Revision	1	Shall be set to 01h to indicate the revision in this specification.
TCGCE-3	2	Event Header Length (EHL)	1	Shall be set to 15h to indicate that 21 bytes of event header information follow and the total event header length (starting from byte 0) is 24 bytes.
TCGCE-4	3	Event Header Additional Information (EHAI)	1	Bits 7:2 are reserved and shall be cleared to 0h. Bits 1:0 shall be set to 01b if the event is associated with an NVM subsystem port and shall be set to 10b if the event is associated with a Management Endpoint.
TCGCE-5	5:4	Controller Identifier	2	Shall contain the NVMe controller identifier for the controller that is associated with the TCG activity.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCGCE-6	13:06	Event Timestamp	8	Shall indicate the Timestamp of when the TCG activity was completed. This field contains an 8-byte data structure whose format matches the data structure returned by a Get Features command that specifies the Timestamp feature (Feature Identifier 0Eh).
TCGCE-7	15:14	Port Identifier (PELPID)	2	Shall contain a Port Identifier as specified in the definition of this field in the NVMe Base Specification.
TCGCE-8	19:16	Reserved	4	Shall be cleared to zero.
TCGCE-9	21:20	Vendor Specific Information Length	2	Shall be set to 04h to align the start of the Vendor Specific Event Data to an 8-byte multiple from the start of this data structure.
TCGCE-10	23:22	Event Length	2	Shall be set to 30h to indicate that 48 bytes of vendor specific information and vendor specific event data follow the event header.
TCGCE-11	27:24	Vendor Specific Information	4	Shall be cleared to zero for padding to align fields that follow.
TCGCE-12	29:28	Vendor Specific Event Code	2	Shall be set to 0001h.
TCGCE-13	30	Vendor Specific Event Data Type	1	Shall be set to 3h to indicate Binary data.
TCGCE-14	31	UUID Index	1	Shall be set to the UUID Index of the UUID (UUID-1) in the UUID List (NVMe-AD-10).
TCGCE-15	33:32	Vendor Specific Event Data Length (VSEDL)	2	Shall be set to 26h to indicate that the length of the vendor specific event data is 38 bytes.
TCGCE-16	35:34	Reserved	2	Shall be cleared to zero.
TCGCE-17	39:36	TCG Command Count	4	Shall increment every time a TCG command is completed regardless of the result and an entry is recorded per TCGHST-PE-2 and TCGHST-PE-3. This value shall be cleared to zero when the device is shipped from the factory.
TCGCE-18	47:40	Invoking ID	8	The 8-byte UID of the table, object, or SP associated with the event (e.g., for a method, this indicates the entity upon which the method is invoked).

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCGCE-19	55:48	Method ID	8	The 8-byte UID of the method being invoked.
TCGCE-20	57:56	ComID	2	The 2-byte ComID specified by the host in the command.
TCGCE-21	58	ProtocolID	1	The Protocol Id specified by the host in the command.
TCGCE-22	59	Status	1	The TCG status on completion of the command.
TCGCE-23	61:60	Process Time	2	Time in milliseconds it took to process the command.
TCGCE-24	71:62	TCG Activity Specific Context	10	Additional context for a TCG activity. If an activity has no additional context this field shall be cleared to 0h. Any unused bytes shall be cleared to 0h.

4.8.3.1 TCG Activity Specific Context

This defines the additional command specific context for specific TCG activities (TCGCE-24).

4.8.3.1.1 Authenticate Method

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCG-AUTHM-1	69:62	Authority ID	8	The 8-byte UID of the authority on which the authenticate method is being invoked.
TCG-AUTHM-2	70	Tries Count	1	The value of the Tries column of the C_PIN object corresponding to the authority after the authenticate method completes.
TCG-AUTHM-3	71	Reserved	1	Shall be cleared to zero.

4.8.3.1.2 Activate Method

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCG-ACT-1	62	Range Start Range Length Policy	1	The value of the RangeStartRangeLengthPolicy parameter passed to the Activate method.
TCG-ACT-2	71:63	Reserved	9	Shall be cleared to zero.

4.8.3.1.3 Reactivate Method

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCG-RACT-1	62	Range Start Range Length Policy	1	The value of the RangeStartRangeLengthPolicy parameter passed to the Reactivate method.
TCG-RACT-2	71:63	Reserved	9	Shall be cleared to zero.

4.8.3.1.4 Assign Method

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCG-ASGN-1	65:62	Namespace Id	4	The value of the NamespaceID parameter passed to the Assign method.
TCG-ASGN-2	71:66	Reserved	6	Shall be cleared to zero.

4.8.3.1.5 Block SID Authentication Command

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TCG-BSID-1	62	Clear Events	1	The value of the Clear Events parameter passed to the Block SID Authentication command.
TCG-BSID-2	71:63	Reserved	9	Shall be cleared to zero.

4.8.4 DSSD Log Page Requirements

The table below defines the scope for all DSSD specific log pages:

Log Identifier	Scope	Log Page Name	Reference Section
C0h	Controller	SMART / Health Information Extended	4.8.6
C1h	NVM subsystem	Error Recovery	4.8.7
C2h	Obsolete		
C3h	Controller	Latency Monitor	4.8.8
C4h	NVM subsystem	Device Capabilities	4.8.10
C5h	NVM subsystem	Unsupported Requirements	4.8.12
C6h	NVM subsystem	Hardware Component	4.8.13
C7h	NVM subsystem	TCG Configuration	4.8.14
C8h	Reserved for future use		
C9h	NVM subsystem	Telemetry String Log	4.8.15
KEY: • Namespace = The log page contains information about a specific namespace. • Controller = The log page contains information about the controller that is processing the command. • NVM subsystem = The log page contains information about the NVM subsystem.			

4.8.5 SMART / Health Information Requirements

Below are the requirements for the SMART / Health Information (Log Identifier 02h) and SMART / Health Information Extended (Log Identifier C0h):

Requirement ID	Description
SLOG-1	All values in the DSSD log pages defined by this specification shall be persistent across power cycles unless otherwise specified.
SLOG-2	All counters defined by this specification shall be saturating counters (i.e., if the counter reaches the maximum allowable size, it stops incrementing and does NOT roll back to 0) unless otherwise specified.
SLOG-3	All fields in DSSD log pages shall be little endian format.
SLOG-4	A normalized counter defined by this specification, unless otherwise specified, shall be reported as the following: 100% shall represent the number at factory exit. 1% shall represent the minimum amount for the device to be reliable. A value of 0% means the device shall no longer be considered reliable. 100% shall be represented as 64h.
SLOG-5	The device shall support all fields in the SMART / Health Information log page.
SLOG-6	A Get Log Page command for either the SMART /Health Information (Log Identifier 02h) or SMART / Health Information Extended (Log Identifier C0h) shall not require an update of the SMART values other than the temperature values in SLOG-8. It shall be a simple read of the current data and shall not block IO.
SLOG-7	Unless otherwise specified, the device shall update the values of the SMART /Health Information (Log Identifier 02h) and SMART / Health Information Extended (Log Identifier C0h) in the background at least once every ten minutes.
SLOG-8	The composite and raw temperature sensor values shall be updated when the log page is accessed.
SLOG-9	All assert events and controller-initiated log captures will require an associated vendor-specific “Reason Identifier” that uniquely identifies the assert /controller condition.
SLOG-10	Obsolete (see STD-LOG-23).
SLOG-11	The device shall not lose any back up energy source failure information or SMART / Health Information (Log Identifier 02h) critical warnings or SMART / Health Information Extended (Log Identifier C0h) critical warnings including across power cycles/resets.
SLOG-12	The device shall not set a SMART/Health critical warning solely due to the value in the Percentage Used field in the SMART/Health log.

4.8.6 SMART / Health Information Extended (Log Identifier C0h)

This vendor-specific log page C0h shall be 512 bytes with the following functional requirements and field format:

Requirement ID	Byte Address	Field	# of Bytes	Field Description
SMART-1	15:0	Physical Media Units Written	16	Shall contain the number of bytes written to the media; this value includes both user and metadata written to the user and system areas. It shall be possible to use this attribute to calculate the Write Amplification Factor (WAF).

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>SMART-2</td><td>31:16</td><td>Physical Media Units Read</td><td>16</td><td>Shall contain the number of bytes read from the media from both the user and system areas.</td></tr>
<tr><td>SMART-3</td><td>39:32</td><td>Bad User NAND Blocks</td><td>8</td><td>The Raw count specifies the number of user NAND blocks that have been retired for any reason (e.g., program fails, erase fails or other events). The Normalized value is the percentage of user spare blocks still available. The normalized value shall be set to 64h, and the Raw count shall be cleared to zero on factory exit.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>39:38</td><td>Normalized value</td></tr>
<tr><td>37:32</td><td>Raw count</td></tr>
</table>
It should be noted that there are 2 bytes for normalized and 6 bytes for raw count. See SLOG-4 definition above.</td></tr>
<tr><td>SMART-4</td><td>47:40</td><td>Bad System NAND Blocks</td><td>8</td><td>The Raw count specifies the number of system NAND blocks that have been retired for any reason (e.g., program fails, erase fails or other events). The Normalized value is the percentage of system spare blocks still available. The normalized value shall be set to 64h, and the Raw count shall be cleared to zero on factory exit. It should be noted that there are 2 bytes for normalized and 6 bytes for raw count. See SLOG-4 definition above.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>47:46</td><td>Normalized value</td></tr>
<tr><td>45:40</td><td>Raw count</td></tr>
</table>
A value of FFFF_FFFF_FFFF_FFFFh indicates that the Bad User NAND block count field above represents all blocks on the device and Bad System NAND block count field is invalid.</td></tr>
<tr><td>SMART-5</td><td>55:48</td><td>XOR Recovery Count</td><td>8</td><td>The total number of times XOR was invoked to recover data in NAND. This shall cover all reads from NAND. Data recovery may have succeeded or failed. This shall be cleared to zero on factory exit.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>SMART-6</td><td>63:56</td><td>Uncorrectable Read Error Count</td><td>8</td><td>Total count of NAND reads that were not correctable by read retries, all levels of ECC, or XOR. This shall be a count of the number of times data recovery fails and an uncorrectable read error is returned to the host.</td></tr>
<tr><td>SMART-7</td><td>71:64</td><td>Soft ECC Error Count</td><td>8</td><td>Total count of NAND reads that were not correctable by first level ECC and requires invoking additional recovery. This shall cover all NAND read accesses. Data recovery may have succeeded or failed. If the device has more than one recovery level, then this counter only increments once per recovery action.</td></tr>
<tr><td>SMART-8</td><td>79:72</td><td>End to End Correction Counts</td><td>8</td><td>The Detected Error count is a count of the detected errors by the end-to-end error protection mechanisms which includes DRAM, SRAM, or other storage element ECC/CRC protection mechanism (not NAND ECC). The Detected Error counter shall increment on all detected end-to-end errors regardless of if the error is corrected or uncorrectable.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>79:76</td><td>Corrected Errors</td></tr>
<tr><td>75:72</td><td>Detected Errors</td></tr>
</table></td></tr>
<tr><td>SMART-9</td><td>80</td><td>System Data % Used</td><td>1</td><td>A normalized cumulative count of the number of erase cycles per block since leaving the factory for the system (firmware and metadata) area. Starts at 0 and increments. 100 indicates that the estimated endurance has been consumed. Value may exceed 100 up to 255. This count shall increment regardless of what the backing media of the blocks are (e.g., SLC and TLC). If system data is split between media types, then this shall report the worst-case count so that the device wear out is clearly understood. This counter has a different behavior than the normalized counter definition in SLOG-4, 100% (64h) represents the device many no longer function reliably as the max erase cycles has been hit.</td></tr>
<tr><td>SMART-10</td><td>87:81</td><td>Refresh Counts</td><td>7</td><td>This is a count of the number of blocks that have been re-allocated to maintain data integrity. This counter does not include creating free space due to garbage collection or block reallocation due to wear leveling.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>SMART-11</td><td>95:88</td><td>User Data Erase Counts</td><td>8</td><td>The maximum count and the minimum erase count across the user NAND blocks in the device. The host shall not be able to reset this counter. It should be noted there are 4 bytes for the maximum and 4 bytes for the minimum. The Minimum User Data Erase Count shall not include bad blocks. If a block goes bad, any subsequent attempts to recover the block shall not increment the Maximum User Data Erase Count field.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>95:92</td><td>Minimum User Data Erase Count</td></tr>
<tr><td>91:88</td><td>Maximum User Data Erase Count</td></tr>
</table></td></tr>
<tr><td>SMART-12</td><td>97:96</td><td>Thermal Throttling Status and Count</td><td>2</td><td>The current status of thermal throttling (throttled or not throttled) and a count of the number of thermal throttling events. Note that there is 1 byte for the current status and 1 byte for the count. This shall be cleared to zero on factory exit.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>97</td><td>Current Throttling Status:
• 00h = unthrottled
• 01h = first level throttle
• 02h = 2nd level throttle
• 03h = 3rd level throttle
• 04h – FFh = Reserved</td></tr>
<tr><td>96</td><td>Number of thermal throttling events</td></tr>
</table></td></tr>
<tr><td>SMART-13</td><td>103:98</td><td>DSSD Specification Version</td><td>6</td><td>Version of the DSSD Specification that this device conforms to.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>103</td><td>Major Version Field. Shall be 02h.</td></tr>
<tr><td>102:101</td><td>Minor Version Field. Shall be 0005h.</td></tr>
<tr><td>100:99</td><td>Point Version Field. Shall be 0000h.</td></tr>
<tr><td>98</td><td>Errata Version Field. Shall be 00h.</td></tr>
</table></td></tr>
<tr><td>SMART-14</td><td>111:104</td><td>PCIe Correctable Error Count</td><td>8</td><td>Summation counter of all PCIe correctable errors (bad TLP, bad DLLP, receiver error, replay timeouts, replay rollovers). These counts shall only increment during run time. They shall not increment during</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				training or power fail. This shall be cleared to zero on factory exit.
SMART-15	115:112	Incomplete Shutdowns	4	A count of the number of shutdowns that have occurred that did not completely flush all required user data and metadata to non-volatile memory for any reason. This shall be cleared to zero on factory exit.
SMART-16	119:116	Reserved	4	Shall be cleared to zero.
SMART-17	120	% Free Blocks	1	This field is the normalized counter of the number of NAND blocks which are currently erased and ready to be written with user data, out of the total number of NAND blocks in the spare area. i.e., this is the percentage of erased NAND blocks in spare area, where the spare area is defined as the NAND blocks that contain no valid user data, no valid system data, and are: • waiting to be erased for use to write user data; or • already erased and ready to be used to write user data. i.e., % Free Blocks = A / (A + B) where A is the number of currently erased NAND blocks and B is the number of NAND blocks waiting to be erased.
SMART-18	127:121	Reserved	7	Shall be cleared to zero.
SMART-19	129:128	Capacitor Health	2	This field is an indicator of the capacitor health and represents the capacitor holdup energy margin during operation. If no PLP protection is present a value of FFFFh shall be reported. 100% represents the passing hold up energy threshold when a device leaves manufacturing. Thus, a device will typically report greater than 100% in this field after leaving manufacturing at the beginning of life. 1% is the minimum hold up energy required to conduct a proper shutdown reliably. A value of 0% may or may not result in a device failing to shutdown properly. This value shall never go negative. Zero is the minimum.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				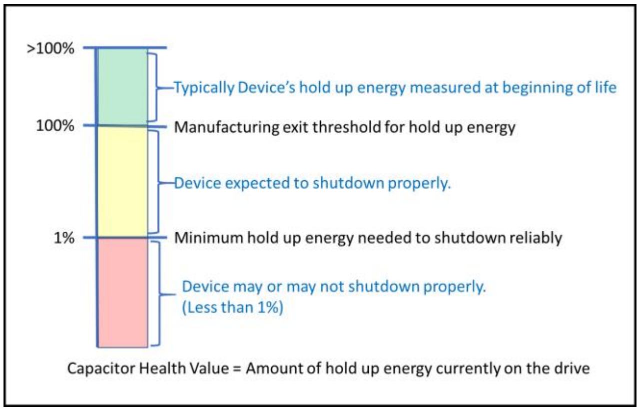 >100% 100% 1% Typically Device's hold up energy measured at beginning of life Manufacturing exit threshold for hold up energy Device expected to shutdown properly. Minimum hold up energy needed to shutdown reliably Device may or may not shutdown properly. (Less than 1%) Capacitor Health Value = Amount of hold up energy currently on the drive
SMART-20	130	NVMe Base Errata Version	1	The device shall report the most recent NVMe Base Specification Errata revision (a, b, c, etc.) in ASCII that the device has implemented backwards incompatible changes from. If the device does not implement any backwards incompatible changes due to errata, this field shall be cleared to zero.
SMART-32	131	NVMe Command Set Errata Version	1	The device shall report the most recent NVMe NVM Command Set Specification Errata revision (a, b, c, etc.) in ASCII that the device has implemented backwards incompatible changes from. If the device does not implement any backwards incompatible changes due to errata, this field shall be cleared to zero.
SMART-30	135:132	Reserved	5	Shall be cleared to zero.
SMART-21	143:136	Unaligned I/O	8	This is a count of the number of write Ios performed by the device that are not aligned to the indirection unit size (IU) of the device. Alignment indicates only the start of each IO. The length does not affect this count. This counter shall reset on power cycle. This counter shall not wrap. This shall be cleared to zero on factory exit.
SMART-22	151:144	Security Version Number	8	This is the Security Version Number of the currently running firmware image. The value of this field shall be 0h for the initial release of firmware. The supplier shall increment this field by one any time a firmware includes a fix for a security issue or critical firmware fix that customer agrees rollback prevention is required.
SMART-23	159:152	Total NUSE	8	Total Namespace Utilization. For a device with a single Namespace, this shall be a copy of the Namespace Utilization field defined in the Identify Namespace Data Structure bytes 23:16. For a device

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				with multiple Namespaces, this shall reflect the total utilization based on all the Namespaces.
SMART-24	175:160	PLP Start Count	16	This is a count of the number of times the device has initiated its power loss protection process due to supply voltage drop. This counter shall be incremented on the initial detection of the power loss condition. This does not include PLP health check operations.
SMART-25	191:176	Endurance Estimate	16	This field is an estimate of the total number of data bytes that may be written to the device over its lifetime assuming a write amplification of 1. (i.e., no increase in the number of write operations performed by the device beyond the number of write operations requested by a host). This value shall be equivalent to the Endurance Estimate field in the Endurance Group Log (Log Identifier 09h).
SMART-29	199:192	PCIe Link Retraining Count	8	This is a count of the number of PCIe Link Retraining events (including link speed and link width changes). This count shall only increment during runtime. It shall not increment during training due to resets or power cycles. This shall be cleared to zero on factory exit.
SMART-31	207:200	Power State Change Count	8	Summation counter of the number of NVMe Power State changes whether host or device initiated, including NVMe Power State changes caused by DSSD Power State changes (see Section 4.8.11 DSSD Power State Requirements). This count shall only increment during run time. This shall be cleared to zero on factory exit.
SMART-33	215:208	Lowest Permitted Firmware Revision	8	This field indicates the value of the Firmware Revision field of the lowest firmware revision that is permitted to be rolled back to. If the currently active firmware image permits rollback to all prior firmware revision, then this field shall be cleared to 0h. If the currently active firmware image does not permit rollback to all prior firmware revisions, then this field shall indicate the value of the Firmware Revision field of the lowest firmware revision that has a Firmware Security Version equal to the Firmware Security Version of the currently active firmware image.
SMART-26	493:216	Reserved	278	Shall be cleared to zero.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
SMART-27	495:494	Log Page Version	2	This indicates the version of the mapping this log page uses. Shall be set to 0004h.
SMART-28	511:496	Log Page GUID	16	Shall be set to AFD514C97C6F4F9CA4f2BFEA2810AFC5h.

4.8.7 Error Recovery (Log Identifier C1h)

Below are the requirements for the Error Recovery (Log Identifier C1h) AEN requirements:

Requirement ID	Description
EREC-AEN-1	If the device uses Panic AEN to report a Panic condition and the device encounters a Panic condition, it shall complete an outstanding Asynchronous Event Request (AER) command with an Asynchronous Event Notification (AEN) with Completion Queue Entry Dword 0 set as follows: • The Log Page Identifier field shall be set to C1h. • The Asynchronous Event Information field shall be cleared to zero. • The Asynchronous Event Type field shall be set to 111b (Vendor Specific).

4.8.7.1 Error Recovery Theory of Operations

Error Recovery Workflow using Asynchronous Event Notifications.

Requirement ID	Description
EREC-FUNC-1	The device shall make every attempt to transparently recover from errors and panics without invoking any of the Error Recovery workflows described below.

Panic Workflow – AEN with in-band recovery

#	Task	Description
		Pre-Failure Parameter Exchange
1	Boot time controller initialization	During boot time controller initialization, the host reads the Error Recovery (Log Identifier C1h) and caches the Panic Reset Wait Time and the Panic Reset Action bitfield. The host also sends an Asynchronous Event Request (AER) command to device and waits for the device to complete the command when there is an asynchronous event.
		Re-Establishing Communication When Failure Occurs
2	Device hits panic condition	The device detects panic condition
3	Device saves panic related data	The device saves the following info: • Panic ID • Device Recovery Action • Debug data for Telemetry log
4	Device generates AEN	The device generates AEN to let the host know that a panic was detected it then starts its panic handling workflow. The device is not

#	Task	Description
		required to handle any NVMe command(s) sent via the Admin Submission Queue after sending the AEN. It is strongly recommended that the device continues to service commands sent via the out-of-band mechanism, if possible. The AEN Completion Queue Entry Dword 0 is filled in per EREC-AEN-1.
5	Host AEN handling	The host services the AEN and logs an event for the panic condition.
6	Host waits for device panic handling to complete	The host waits for Panic Reset Wait Time to allow the device to finish its panic handling workflow.
Extracting Debug Info After Failure and Recovery		
7	Host recovers device	The host tries one or more of the resets specified in the Panic Reset Action bitfield to attempt to bring the device back to a state where it can service NVMe commands.
8	Host initiates controller initialization	The host starts controller initialization. If the controller fails to initialize, the host will attempt the next Panic Reset Action as specified in the Panic Reset Action bitfield.
9	Host reads 0xC1 log page	The host reads the Error Recovery (Log Identifier C1h) to retrieve the following info: • Panic ID • Device Recovery Action A non-zero Panic ID indicates the device is in panic mode. The device fails IO commands while it is in panic mode and cannot safely complete commands: • Status Code (SC) = 06h (Internal Error) • Status Code Type (SCT) = 00h (Generic Command Status)
10	Host retrieves Controller-Initiated Telemetry Log if device in panic mode	Host retrieves Telemetry Controller-Initiated (Log Identifier 08h) for later use in debugging.
11	Device Recovery Action	The recommended Device Recovery Action is taken. This is done during the offline recovery workflow.

Panic Workflow – CFS with in-band recovery

#	Task	Description
Pre-Failure Parameter Exchange		
1	Boot time controller initialization	During boot time controller initialization, driver reads 0xC1 log page and caches Panic Reset Wait Time and Panic Reset Action. Driver also saves info that Panic CFS Supported is set.
Re-Establishing Communication When Failure Occurs		
2	Device hits panic condition	Device detects panic condition

#	Task	Description
3	Device saves panic related data	Device saves the following info: • Panic ID • Device Recovery Action Debug data for controller initiated log
4	Device asserts CFS	Device asserts CFS to let host know that a panic was detected and starts panic handling workflow. Device is not expected to be able to handle any NVMe command(s) sent via the Admin Submission Queue after asserting CFS. It is strongly recommended that the device continues to service commands sent via the out-of-band mechanism, if possible.
5	Driver detects CFS after command timeout	On first command timeout after CFS bit is set, driver will check and detect CFS bit is set.
6	Driver waits for device panic handling to complete	The driver waits for Panic Reset Wait Time to allow the device to finish its panic handling workflow.
	Extracting Debug Info After Failure and Recovery	
7	Driver initiates a NVMe Controller Reset	The driver initiates a NVMe Controller Reset as recommended in NVMe Spec for CFS.
8	Driver resets device if NVMe Controller Reset fails	If prior NVMe Controller Reset fails, driver tries one or more resets specified in Panic Reset Action to attempt to bring device back to a state that can service NVMe command(s).
9	Driver initiates controller initialization	Driver starts controller initialization. If controller fails to initialize, driver may attempt a different Panic Reset Action.
10	Driver reads 0xC1 log page	Driver reads 0xC1 log page to retrieve the following info: • Panic ID • Device Recovery Action A non-zero Panic ID indicates device is in panic mode. Fail IO Commands while drive is in panic mode (unless drive is brick) and cannot safely complete commands: • Status Code (SC) = 0x06 (Internal Error) • Status Code Type (SCT) = 0x0 (Generic Command Status)
11	Driver retrieves Controller-Initiated Telemetry Log if device in panic mode	Driver retrieves Controller-Initiated Telemetry Log and uploads the log for sharing with IHV
12	Device Recovery Action	The recommended Device Recovery Action is taken. This is done during the offline recovery workflow.

This vendor-specific log page, C1h shall be 512 bytes with the following functional requirements and field format:

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EREC -1</td><td>1:0</td><td>Panic Reset Wait Time</td><td>2</td><td>The amount of time the host should wait for the device panic workflow to complete in msec.</td></tr>
<tr><td>EREC-2</td><td>2</td><td>Panic Reset Action</td><td>1</td><td>Bit field indicating potential reset actions that may need to be taken. If no reset action is needed, do not set any of the bits. More than 1 bit can be set, and it is up to the host to decide the sequence of action(s) to take.
Instead of asserting and requiring the host to perform an NVMe Controller Reset to clear the assert condition, the device should reset itself using a methodology similar to what is used to perform a firmware activation without reset (e.g., pause command processing, set the Processing Paused bit to 1b, save runtime state, re-initialize firmware, restore runtime state, resuming processing, clear the Processing Paused bit to 0b). The device shall log this event in the telemetry logs and include a vendor specific panic code in the Error ID (see TELRI-1) whether or not it can autonomously recover.
If autonomous recovery is not possible, then the assert condition shall be reported to the host (see EREC-5).
The preferred order of device advertised recovery is bit 0, 1, 2, 5, 3 then 4 which is based on least impactful to most impactful from a host perspective.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:6</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>5</td><td>PCIe Conventional Hot Reset.</td></tr>
<tr><td>4</td><td>Main Power Cycle.</td></tr>
<tr><td>3</td><td>PERST#.</td></tr>
<tr><td>2</td><td>PCIe Function Level Reset.</td></tr>
<tr><td>1</td><td>NVM Subsystem Reset.</td></tr>
<tr><td>0</td><td>NVMe Controller Reset.</td></tr>
</table></td></tr>
<tr><td>EREC-3</td><td>3</td><td>Device Recovery Action 1</td><td>1</td><td>The recovery action to take for handling a device panic condition. Value is dependent on the panic condition. The device shall set bit 0 to 1b if data integrity is intact. The device shall set bit 6 to 1b if the data integrity of some of the LBAs is compromised. The device shall only set bit 1, bit 2 or bit 5 to 1b if the data integrity of all of the LBAs is compromised. The device shall only set bit 3 to 1b or</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>bit 4 to 1b if there is a persistent hardware failure. The device shall not set more than one bit to one in this field.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>6</td><td>User data loss. The device shall populate LBA Status Information (Log Identifier 0Eh) with the lost/corrupted LBAs.</td></tr>
<tr><td>5</td><td>Sanitize Required.</td></tr>
<tr><td>4</td><td>Device Replacement Required.</td></tr>
<tr><td>3</td><td>Vendor Analysis Required.</td></tr>
<tr><td>2</td><td>Vendor Specific Command Required.</td></tr>
<tr><td>1</td><td>Format NVM Required (Any SES value, any supported combination of other parameters – PIL, PI, MSET, LBAF).</td></tr>
<tr><td>0</td><td>No Action Required.</td></tr>
</table></td></tr>
<tr><td>EREC-4</td><td>11:4</td><td>Panic ID</td><td>8</td><td>A numeric value that uniquely identifies the exact panic condition encountered. A value of zero indicates no panic.
The following Panic ID values are reserved for Host defined fault codes for known panic conditions:
• 00000000_00000000h – 00000000_0000FFFFh
<table>
<tr><th>Byte</th><th>Byte Description</th></tr>
<tr><td>11:4</td><td>Panic ID definition:
• 00000000_00000001h – Panic caused by flush failures or data loss during power loss handling.</td></tr>
</table></td></tr>
<tr><td>EREC-5</td><td>15:12</td><td>Device Capabilities</td><td>4</td><td>Field to indicate device capabilities.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>31:2</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>1</td><td>Panic CFS Supported: If set, indicates device supports using CFS to notify host of a panic condition*.</td></tr>
<tr><td>0</td><td>Panic AEN Supported: If set, indicates device supports using AEN to notify host of a panic condition*.</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>*Note: It is valid for a device to indicate support for both Panic AEN Supported and Panic Controller Fatal Status Supported. If the device supports both, the device shall only use one of these in-band panic notification mechanisms when reporting a given panic event as configured in the DSSD Asynchronous Event Configuration Feature (see SDAEC-14).</td></tr>
<tr><td>EREC-6</td><td>16</td><td>Vendor Specific Recovery Opcode</td><td>1</td><td>Vendor specific command opcode to recover device from panic condition. Only valid when bit 2 of the Device Recovery Action 1 field value is set to 1b. When bit 2 of the Device Recovery Action 1 field value is cleared to 0b, this field shall be cleared to zero.</td></tr>
<tr><td>EREC-7</td><td>19:17</td><td>Reserved</td><td>3</td><td>Shall be cleared to zero.</td></tr>
<tr><td>EREC-8</td><td>23:20</td><td>Vendor Specific Command CDW12</td><td>4</td><td>CDW12 value for the Vendor Specific command to recover device from panic condition. Only valid when bit 2 of the Device Recovery Action 1 field is set to 1b. When bit 2 of the Device Recovery Action 1 field value is cleared to 0b, this field shall be cleared to zero.</td></tr>
<tr><td>EREC-9</td><td>27:24</td><td>Vendor Specific Command CDW13</td><td>4</td><td>CDW13 value for the Vendor Specific command to recover device from panic condition. Only valid when bit 2 of the Device Recovery Action 1 field value is set to 1b. When Device Recovery Action 1 field value is cleared to 0b, this field shall be cleared to zero.</td></tr>
<tr><td>EREC-13</td><td>28</td><td>Vendor Specific Command Timeout</td><td>1</td><td>The amount of time the host should wait for the device to complete the recovery command in seconds.</td></tr>
<tr><td>EREC-14</td><td>29</td><td>Device Recovery Action 2</td><td>1</td><td>Bit field indicating potential post reset actions that may need to be taken. If no reset action is needed, do not set any of the bits. More than 1 bit can be set, and it is up to the host to decide the sequence of action(s) to take. Use Bit 0 if possible. If Bit 0 is not possible then use Bit 1, etc.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:6</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>5</td><td>PCIe Conventional Hot Reset.</td></tr>
<tr><td>4</td><td>Main Power Cycle.</td></tr>
<tr><td>3</td><td>PERST#.</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td>2</td><td>PCIe Function Level Reset.</td></tr><tr><td>1</td><td>NVM Subsystem Reset.</td></tr><tr><td>0</td><td>NVMe Controller Reset.</td></tr></table></td></tr>
<tr><td>EREC-15</td><td>30</td><td>Device Recovery Action 2 Timeout</td><td>1</td><td>The amount of time the host should wait for the device to complete the Device Recovery Action 2 in seconds (overrides CAP.TO value).</td></tr>
<tr><td>EREC-16</td><td>31</td><td>Panic Count</td><td>1</td><td>The number of Panics the device has encountered during its lifetime. This is a saturating counter.</td></tr>
<tr><td>EREC-17</td><td>63:32</td><td>Previous Panic IDs</td><td>32</td><td>This set of fields reports the four most recent Panic IDs not including the ID of the current panic.<table><tr><th>Byte</th><th>Byte Description</th></tr><tr><td>39:32</td><td>Panic ID N-1</td></tr><tr><td>47:40</td><td>Panic ID N-2</td></tr><tr><td>55:48</td><td>Panic ID N-3</td></tr><tr><td>63:56</td><td>Panic ID N-4</td></tr></table></td></tr>
<tr><td>EREC-10</td><td>493:64</td><td>Reserved</td><td>463</td><td>Shall be cleared to zero.</td></tr>
<tr><td>EREC-11</td><td>495:494</td><td>Log Page Version</td><td>2</td><td>This indicates the version of the mapping this log page uses. Shall be set to 0003h.</td></tr>
<tr><td>EREC-12</td><td>511:496</td><td>Log Page GUID</td><td>16</td><td>Shall be set to 5A1983BA3DFD4DABAE3430FE2131D944h.</td></tr>
</table>

4.8.8 Latency Monitor Log and Feature Set Requirements

The following are requirements for the Latency Monitor Log and Feature Set. For more information about the Latency Monitoring Feature Set (see Section 22 Latency Monitoring Feature Set Theory of Operation).

Requirement ID	Description
LMLOG-1	All values in the Latency Monitor Log Page shall be persistent across power cycles and resets unless otherwise specified.
LMLOG-2	All counters shall be saturating counters (i.e., if the counter reaches the maximum allowable size, it stops incrementing and does NOT roll back to 0).
LMLOG-3	All values in the Latency Monitor Log shall be little endian format.
LMLOG-4	A read of the Latency Monitor Log shall be a simple read of the active data and shall not block IO.
LMLOG-5	Data in the Latency Monitor Log which is read by the host shall be no more than 10 minutes old.
LMLOG-6	When configuring the Latency Monitoring Feature with Set Features the Active Buckets and Static Buckets shall be reset.

Requirement ID	Description
LMLOG-7	When powering on, the counters shall run based on the previously configured values. If the values have never been configured, they shall run based on the default values.
LMLOG-8	The Latency Stamp shall be based on the NVMe timestamp set by the last Timestamp (Feature 0Eh) Set Features command if any. If a NVMe Timestamp (Feature 0Eh) Set Features command has not been received, then the Latency Stamp shall be based on the device power on hours.
LMLOG-9	When the device provides a Latency Stamp of latency outliers, the Latency Stamp shall be based on command completion.
LMLOG-10	When generating a Latency Monitoring Log, the latency shall be no greater than the latency associated with generating a Telemetry log.
LMLOG-11	When configuring this feature, the thresholds shall always be configured such that Active Threshold A < Active Threshold B < Active Threshold C < Active Threshold D. If the host attempts to configure the device in such a way that violates the above rules, the device shall return Invalid Field in Command.
LMLOG-12	The device shall support the Latency Monitoring Feature Set.
LMLOG-13	The Latency Monitor Log page shall be 512 bytes.
LMLOG-14	Executing a Set Feature command for Latency Monitor (Feature Identifier C5h) shall reset all the contents in the Active/Static Buckets.

4.8.9 Latency Monitor (Log Identifier C3h)

This vendor-specific log page, C3h shall be 512 bytes with the following functional requirements and field format (see Section 22 Latency Monitoring Feature Set Theory of Operation for additional details):

Requirement ID	Byte Address	Field	# of Bytes	Field Description	
				Bit	Bit Description
LMDATA-1	0	Latency Monitor Feature Status	1	7:3	Reserved. Shall be cleared to zero.
				2	Active Measured Latency Supported. When set to 1b the Active Measured Latency field is supported and shall be populated based on the Active Latency Configuration settings. When cleared to 0b the Active Measured Latency is not supported.
				1	Active Latency Configuration/Active Latency Mode 1 Supported. When set to 1b the device shall support the Active Latency Configuration with the Active Latency Mode = 1b. When cleared to 0b the device does not support Active Latency Mode =1b.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				0 — Latency Monitoring Feature Enabled. This is a global feature enable. When set to 1b the Latency Monitoring Feature for this device is enabled. When cleared to 0b all features in the Latency Monitoring Log page are disabled for this device and can be ignored. When cleared to 0b the other fields in this log page are not power loss safe and may be lost. The default value shall be 07h.
LMDATA-2	1	Reserved	1	Shall be cleared to zero.
LMDATA-3	3:2	Active Bucket Timer	2	The Active Bucket Timer is in 5-minute increments. Thus, a value of 0001h is 5 minutes and a value of 0002h is 10 minutes. This represents the amount of time the Active Buckets have been accumulating data. The Active Bucket Timer will saturate at FFFFh. When the Active Bucket Timer reaches the Active Bucket Timer Threshold then the data specified in Section 22 Latency Monitoring Feature Set Theory of Operation is moved into the Static Buckets, the Active Bucket Timer is cleared to 0000h and restarts counting. If the Active Bucket Timer is running and there is a power cycle the Active Bucket Timer value from before the power cycle shall be restored into the Active Bucket Timer and the Active Bucket Timer shall continue when the device is powered on.
LMDATA-4	5:4	Active Bucket Timer Threshold	2	Active Bucket Timer Threshold is the threshold used to compare with the Active Bucket Timer. When cleared to 0000h this threshold is not used, the Active Bucket Timer will saturate, and the Static Buckets will not be loaded. This threshold is in 5-minute increments. The factory default value of the Active Bucket Timer Threshold shall be set to 07E0h.
LMDATA-5	6	Active Threshold A	1	This defines Active Threshold A. This is in 5ms increments. A value of 00h represents 5ms. A value of FFh represents 1.280 seconds. The factory default is 05h.
LMDATA-6	7	Active Threshold B	1	This is in 5ms increments. The factory default is 13h.

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>LMDATA-7</td><td>8</td><td>Active Threshold C</td><td>1</td><td>This is in 5ms increments. The factory default is 1Eh.</td></tr>
<tr><td>LMDATA-8</td><td>9</td><td>Active Threshold D</td><td>1</td><td>This is in 5ms increments. The factory default is 2Eh.</td></tr>
<tr><td>LMDATA-9</td><td>11:10</td><td>Active Latency Configuration</td><td>2</td><td>This configures how both the Active Latency Stamp, and the Active Measured Latency Fields are updated on a per I/O command (Read, Write, Deallocate) counter basis.
When the Active Latency Mode is cleared to 0b the Active Latency Stamp and the Active Measured Latency will trigger and update the first time the associated command counter increments. Once this trigger happens the fields shall not be updated until the Active Latency Stamp and Active Measured Latency fields are reset based on the Active Bucket Timer expiring.
When the Active Latency Mode is set to 1b the Active Latency Stamp and the Active Measured Latency fields shall update to show the largest measured latency based on the associated command counter.
<table>
<tr><th>Bit</th><th>Bucket</th><th>Counter</th><th>Bit Description</th></tr>
<tr><td>15:12</td><td>N/A</td><td>N/A</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>11</td><td>3</td><td>De-allocate/TRIM</td><td>Active Latency Mode [11]</td></tr>
<tr><td>10</td><td>3</td><td>Write</td><td>Active Latency Mode [10]</td></tr>
<tr><td>9</td><td>3</td><td>Read</td><td>Active Latency Mode [9]</td></tr>
<tr><td>8</td><td>2</td><td>De-allocate/TRIM</td><td>Active Latency Mode [8]</td></tr>
<tr><td>7</td><td>2</td><td>Write</td><td>Active Latency Mode [7]</td></tr>
<tr><td>6</td><td>2</td><td>Read</td><td>Active Latency Mode [6]</td></tr>
<tr><td>5</td><td>1</td><td>De-allocate/TRIM</td><td>Active Latency Mode [5]</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>4</td><td>1</td><td>Write</td><td>Active Latency Mode [4]</td></tr>
<tr><td>3</td><td>1</td><td>Read</td><td>Active Latency Mode [3]</td></tr>
<tr><td>2</td><td>0</td><td>De-allocate/TRIM</td><td>Active Latency Mode [2]</td></tr>
<tr><td>1</td><td>0</td><td>Write</td><td>Active Latency Mode [1]</td></tr>
<tr><td>0</td><td>0</td><td>Read</td><td>Active Latency Mode [0]</td></tr>
</table>
The default value shall be 0FFFh.</td></tr>
<tr><td>LMDATA-10</td><td>12</td><td>Active Latency Minimum Window</td><td>1</td><td>This is the minimum number of 100 millisecond increments between Latency Events for a single Active Latency Stamp and Active Measured Latency. When cleared to 00h this feature is disabled. This count is in 100 millisecond increments, thus a value of 01h is 100 milliseconds and 02h is 200 milliseconds. Once a Latency Stamp/Measured Latency is updated if the Active Latency Minimum Window time has not expired and an event that is configurated to generate a Latency Stamp/Measured Latency occurs the Latency Stamp/Measured Latency will not be recorded. The default value of this field is 0Ah.</td></tr>
<tr><td>LMDATA-11</td><td>31:13</td><td>Reserved</td><td>19</td><td>Shall be cleared to zero.</td></tr>
<tr><td>LMDATA-12</td><td>47:32</td><td>Active Bucket Counter 0</td><td>16</td><td><table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>47:44</td><td>Read Command Counter.</td></tr>
<tr><td>43:40</td><td>Write Command Counter.</td></tr>
<tr><td>39:36</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>35:32</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>LMDATA-13</td><td>63:48</td><td>Active Bucket Counter 1</td><td>16</td><td><table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>63:60</td><td>Read Command Counter.</td></tr>
<tr><td>59:56</td><td>Write Command Counter.</td></tr>
<tr><td>55:52</td><td>De-Allocate/TRIM Command Counter.</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td>51:48</td><td>Reserved. Shall be cleared to zero.</td></tr></table></td></tr>
<tr><td>LMDATA-14</td><td>79:64</td><td>Active Bucket Counter 2</td><td>16</td><td><table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>79:76</td><td>Read Command Counter.</td></tr>
<tr><td>75:72</td><td>Write Command Counter.</td></tr>
<tr><td>71:68</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>67:64</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>LMDATA-15</td><td>95:80</td><td>Active Bucket Counter 3</td><td>16</td><td><table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>95:92</td><td>Read Command Counter.</td></tr>
<tr><td>91:88</td><td>Write Command Counter.</td></tr>
<tr><td>87:84</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>83:80</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>LMDATA-16</td><td>191:96</td><td>Active Latency Stamp</td><td>96</td><td>This field contains a Timestamp for when a latency event occurred for each counter. A value of FFFFFFFF_FFFFFFFFh means the Latency Stamp is not valid. The Active Latency Stamp uses the data format for Timestamp as defined in NVMe.<table>
<tr><th>Byte Address</th><th>Bucket</th><th>Counter</th><th>Byte Description</th></tr>
<tr><td>191:184</td><td>0</td><td>Read</td><td>Active Latency Stamp 0</td></tr>
<tr><td>183:176</td><td>0</td><td>Write</td><td>Active Latency Stamp 1</td></tr>
<tr><td>175:168</td><td>0</td><td>De-allocate/ TRIM</td><td>Active Latency Stamp 2</td></tr>
<tr><td>167:160</td><td>1</td><td>Read</td><td>Active Latency Stamp 3</td></tr>
<tr><td>159:152</td><td>1</td><td>Write</td><td>Active Latency Stamp 4</td></tr>
<tr><td>151:144</td><td>1</td><td>De-allocate/ TRIM</td><td>Active Latency Stamp 5</td></tr>
<tr><td>143:136</td><td>2</td><td>Read</td><td>Active Latency Stamp 6</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>135:128</td><td>2</td><td>Write</td><td>Active Latency Stamp 7</td></tr>
<tr><td>127:120</td><td>2</td><td>De-allocate/ TRIM</td><td>Active Latency Stamp 8</td></tr>
<tr><td>119:112</td><td>3</td><td>Read</td><td>Active Latency Stamp 9</td></tr>
<tr><td>111:104</td><td>3</td><td>Write</td><td>Active Latency Stamp 10</td></tr>
<tr><td>103:96</td><td>3</td><td>De-allocate/ TRIM</td><td>Active Latency Stamp 11</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-17</td><td>215:192</td><td>Active Measured Latency</td><td>24</td><td>This is the measured latency that caused the counter to increment. A value of 0000h means this field is invalid. A value of 0001h represents 1 millisecond. A value of 0002h represents 2 milliseconds.
<table>
<tr><th>Byte Address</th><th>Bucket</th><th>Counter</th><th>Byte Description</th></tr>
<tr><td>215:214</td><td>0</td><td>Read</td><td>Active Measured Latency 0</td></tr>
<tr><td>213:212</td><td>0</td><td>Write</td><td>Active Measured Latency 1</td></tr>
<tr><td>211:210</td><td>0</td><td>De-allocate/ TRIM</td><td>Active Measured Latency 2</td></tr>
<tr><td>209:208</td><td>1</td><td>Read</td><td>Active Measured Latency 3</td></tr>
<tr><td>207:206</td><td>1</td><td>Write</td><td>Active Measured Latency 4</td></tr>
<tr><td>205:204</td><td>1</td><td>De-allocate/ TRIM</td><td>Active Measured Latency 5</td></tr>
<tr><td>203:202</td><td>2</td><td>Read</td><td>Active Measured Latency 6</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>201:200</td><td>2</td><td>Write</td><td>Active Measured Latency 7</td></tr>
<tr><td>199:198</td><td>2</td><td>De-allocate/TRIM</td><td>Active Measured Latency 8</td></tr>
<tr><td>197:196</td><td>3</td><td>Read</td><td>Active Measured Latency 9</td></tr>
<tr><td>195:194</td><td>3</td><td>Write</td><td>Active Measured Latency 10</td></tr>
<tr><td>193:192</td><td>3</td><td>De-allocate/TRIM</td><td>Active Measured Latency 11</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-18</td><td>217:216</td><td>Active Latency Stamp Units</td><td>2</td><td>When bit is set to 1b, the Active Latency Stamp was based on receiving the NVMe Timestamp. When bit is cleared to 0b the Active Latency Stamp was based on power on hours since the NVMe Timestamp was not received.
<table>
<tr><th>Bit</th><th>Bucket</th><th>Counter</th><th>Bit Description</th></tr>
<tr><td>15:12</td><td>N/A</td><td>N/A</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>11</td><td>3</td><td>De-allocate/TRIM</td><td>Active Latency Stamp Unit [11]</td></tr>
<tr><td>10</td><td>3</td><td>Write</td><td>Active Latency Stamp Unit [10]</td></tr>
<tr><td>9</td><td>3</td><td>Read</td><td>Active Latency Stamp Unit [9]</td></tr>
<tr><td>8</td><td>2</td><td>De-allocate/TRIM</td><td>Active Latency Stamp Unit [8]</td></tr>
<tr><td>7</td><td>2</td><td>Write</td><td>Active Latency Stamp Unit [7]</td></tr>
<tr><td>6</td><td>2</td><td>Read</td><td>Active Latency Stamp Unit [6]</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>5</td><td>1</td><td>De-allocate/TRIM</td><td>Active Latency Stamp Unit [5]</td></tr>
<tr><td>4</td><td>1</td><td>Write</td><td>Active Latency Stamp Unit [4]</td></tr>
<tr><td>3</td><td>1</td><td>Read</td><td>Active Latency Stamp Unit [3]</td></tr>
<tr><td>2</td><td>0</td><td>De-allocate/TRIM</td><td>Active Latency Stamp Unit [2]</td></tr>
<tr><td>1</td><td>0</td><td>Write</td><td>Active Latency Stamp Unit [1]</td></tr>
<tr><td>0</td><td>0</td><td>Read</td><td>Active Latency Stamp Unit [0]</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-19</td><td>239:218</td><td>Reserved</td><td>22</td><td>Shall be cleared to zero.</td></tr>
<tr><td>LMDATA-20</td><td>255:240</td><td>Static Bucket Counter 0</td><td>16</td><td>This is a snapshot of the Active Bucket Counter 0 which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>255:252</td><td>Read Command Counter.</td></tr>
<tr><td>251:248</td><td>Write Command Counter.</td></tr>
<tr><td>247:244</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>243:240</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-21</td><td>271:256</td><td>Static Bucket Counter 1</td><td>16</td><td>This is a snapshot of the Active Bucket Counter 1 which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>271:268</td><td>Read Command Counter.</td></tr>
<tr><td>267:264</td><td>Write Command Counter.</td></tr>
<tr><td>263:260</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>259:256</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-22</td><td>287:272</td><td>Static Bucket Counter 2</td><td>16</td><td>This is a snapshot of the Active Bucket Counter 2 which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>287:284</td><td>Read Command Counter.</td></tr>
<tr><td>283:280</td><td>Write Command Counter.</td></tr>
<tr><td>279:276</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>275:272</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-23</td><td>303:288</td><td>Static Bucket Counter 3</td><td>16</td><td>
This is a snapshot of the Active Bucket Counter 3 which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>303:300</td><td>Read Command Counter.</td></tr>
<tr><td>299:296</td><td>Write Command Counter.</td></tr>
<tr><td>295:292</td><td>De-Allocate/TRIM Command Counter.</td></tr>
<tr><td>291:288</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-24</td><td>399:304</td><td>Static Latency Stamp</td><td>96</td><td>
This is a snapshot of the Active Latency Stamp which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.
This field contains a timestamp for when a latency event occurred for each counter. A value of FFFFFFFF_FFFFFFFFh means the Latency Stamp is not valid. The Static Latency Stamp uses the Timestamp data format as defined in NVMe.
<table>
<tr><th>Byte Address</th><th>Bucket</th><th>Counter</th><th>Byte Description</th></tr>
<tr><td>399:392</td><td>0</td><td>Read</td><td>Static Latency Stamp 0</td></tr>
<tr><td>391:384</td><td>0</td><td>Write</td><td>Static Latency Stamp 1</td></tr>
<tr><td>383:376</td><td>0</td><td>De-allocate/ TRIM</td><td>Static Latency Stamp 2</td></tr>
<tr><td>375:368</td><td>1</td><td>Read</td><td>Static Latency Stamp 3</td></tr>
<tr><td>367:360</td><td>1</td><td>Write</td><td>Static Latency Stamp 4</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>359:352</td><td>1</td><td>De-allocate/ TRIM</td><td>Static Latency Stamp 5</td></tr>
<tr><td>351:344</td><td>2</td><td>Read</td><td>Static Latency Stamp 6</td></tr>
<tr><td>343:336</td><td>2</td><td>Write</td><td>Static Latency Stamp 7</td></tr>
<tr><td>335:328</td><td>2</td><td>De-allocate/ TRIM</td><td>Static Latency Stamp 8</td></tr>
<tr><td>327:320</td><td>3</td><td>Read</td><td>Static Latency Stamp 9</td></tr>
<tr><td>319:312</td><td>3</td><td>Write</td><td>Static Latency Stamp 10</td></tr>
<tr><td>311:304</td><td>3</td><td>De-allocate/ TRIM</td><td>Static Latency Stamp 11</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-25</td><td>423:400</td><td>Static Measured Latency</td><td>24</td><td>
<p>This is a snapshot of the Active Measured Latency which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.</p>
<p>A value of 0000h means this field is invalid. A value of 0001h represents 1 millisecond. A value of 0002h represents 2 milliseconds.</p>
<table>
<tr><th>Byte Address</th><th>Bucket</th><th>Counter</th><th>Byte Description</th></tr>
<tr><td>423:422</td><td>0</td><td>Read</td><td>Static Measured Latency 0</td></tr>
<tr><td>421:420</td><td>0</td><td>Write</td><td>Static Measured Latency 1</td></tr>
<tr><td>419:418</td><td>0</td><td>De-allocate/ TRIM</td><td>Static Measured Latency 2</td></tr>
<tr><td>417:416</td><td>1</td><td>Read</td><td>Static Measured Latency 3</td></tr>
<tr><td>415:414</td><td>1</td><td>Write</td><td>Static Measured Latency 4</td></tr>
<tr><td>413:412</td><td>1</td><td>De-allocate/ TRIM</td><td>Static Measured Latency 5</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>411:410</td><td>2</td><td>Read</td><td>Static Measured Latency 6</td></tr>
<tr><td>409:408</td><td>2</td><td>Write</td><td>Static Measured Latency 7</td></tr>
<tr><td>407:406</td><td>2</td><td>De-allocate/ TRIM</td><td>Static Measured Latency 8</td></tr>
<tr><td>405:404</td><td>3</td><td>Read</td><td>Static Measured Latency 9</td></tr>
<tr><td>403:402</td><td>3</td><td>Write</td><td>Static Measured Latency 10</td></tr>
<tr><td>401:400</td><td>3</td><td>De-allocate/ TRIM</td><td>Static Measured Latency 11</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-26</td><td>425:424</td><td>Static Latency Stamp Units</td><td>2</td><td>This is a snapshot of the Active Latency Stamp Units which is moved to this field when the Active Bucket Timer equals the Active Bucket Timer Threshold.
When bit is set to 1b the Static Latency Stamp was based on receiving the NVMe Timestamp and the offset from this. When bit is cleared to 0b the Static Latency Stamp was based on power on hours since the NVMe Timestamp was not received.
<table>
<tr><th>Bit</th><th>Bucket</th><th>Counter</th><th>Byte Description</th></tr>
<tr><td>15:12</td><td>N/A</td><td>N/A</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>11</td><td>3</td><td>De-allocate/TRIM</td><td>Static Latency Stamp Unit [11]</td></tr>
<tr><td>10</td><td>3</td><td>Write</td><td>Static Latency Stamp Unit [10]</td></tr>
<tr><td>9</td><td>3</td><td>Read</td><td>Static Latency Stamp Unit [9]</td></tr>
<tr><td>8</td><td>2</td><td>De-allocate/TRIM</td><td>Static Latency Stamp Unit [8]</td></tr>
<tr><td>7</td><td>2</td><td>Write</td><td>Static Latency Stamp Unit [7]</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>6</td><td>2</td><td>Read</td><td>Static Latency Stamp Unit [6]</td></tr>
<tr><td>5</td><td>1</td><td>De-allocate/TRIM</td><td>Static Latency Stamp Unit [5]</td></tr>
<tr><td>4</td><td>1</td><td>Write</td><td>Static Latency Stamp Unit [4]</td></tr>
<tr><td>3</td><td>1</td><td>Read</td><td>Static Latency Stamp Unit [3]</td></tr>
<tr><td>2</td><td>0</td><td>De-allocate/TRIM</td><td>Static Latency Stamp Unit [2]</td></tr>
<tr><td>1</td><td>0</td><td>Write</td><td>Static Latency Stamp Unit [1]</td></tr>
<tr><td>0</td><td>0</td><td>Read</td><td>Static Latency Stamp Unit [0]</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-27</td><td>435:426</td><td>Reserved</td><td>10</td><td>Shall be cleared to zero.</td></tr>
<tr><td>LMDATA-37</td><td>447:436</td><td>Latency Monitor Debug Telemetry Log Size</td><td>12</td><td>This is the number of Dwords in the Latency Monitor Debug Log. This value is in Dwords.</td></tr>
<tr><td>LMDATA-28</td><td>449:448</td><td>Debug Log Trigger Enable</td><td>2</td><td>This controls what counters can cause a debug log event to be triggered. When set to 1b the first time the bucket/counter combination is incremented a debug log is triggered. When cleared to 0b a debug log will not be triggered when the bucket/counter combination is incremented.
<table>
<tr><th>Bit</th><th>Default Value</th><th>Bucket</th><th>Counter</th><th>Bit Description</th></tr>
<tr><td>15:12</td><td>N/A</td><td>N/A</td><td>N/A</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>11</td><td>0b</td><td>3</td><td>De-allocate /TRIM</td><td>Active Log Enable [11]</td></tr>
<tr><td>10</td><td>0b</td><td>3</td><td>Write</td><td>Active Log Enable [10]</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>9</td><td>0b</td><td>3</td><td>Read</td><td>Active Log Enable [9]</td></tr>
<tr><td>8</td><td>0b</td><td>2</td><td>De-allocate /TRIM</td><td>Active Log Enable [8]</td></tr>
<tr><td>7</td><td>0b</td><td>2</td><td>Write</td><td>Active Log Enable [7]</td></tr>
<tr><td>6</td><td>0b</td><td>2</td><td>Read</td><td>Active Log Enable [6]</td></tr>
<tr><td>5</td><td>0b</td><td>1</td><td>De-allocate /TRIM</td><td>Active Log Enable [5]</td></tr>
<tr><td>4</td><td>0b</td><td>1</td><td>Write</td><td>Active Log Enable [4]</td></tr>
<tr><td>3</td><td>0b</td><td>1</td><td>Read</td><td>Active Log Enable [3]</td></tr>
<tr><td>2</td><td>0b</td><td>0</td><td>De-allocate /TRIM</td><td>Active Log Enable [2]</td></tr>
<tr><td>1</td><td>0b</td><td>0</td><td>Write</td><td>Active Log Enable [1]</td></tr>
<tr><td>0</td><td>0b</td><td>0</td><td>Read</td><td>Active Log Enable [0]</td></tr>
</table>
</td></tr>
<tr><td>LMDATA-29</td><td>451:450</td><td>Debug Log Measured Latency</td><td>2</td><td>When a debug log is triggered, this is the measured latency for the Latency Stamp that caused the debug log to trigger. A value of 0000h means this field is invalid. A value of 0001h represents 1ms. A value of 0002h represents 2 milliseconds. A value of FFFFh means the Debug Log Measured Latency saturated at the max.</td></tr>
<tr><td>LMDATA-30</td><td>459:452</td><td>Debug Log Latency Stamp</td><td>8</td><td>This is the Latency Stamp associated with the debug log.</td></tr>
<tr><td>LMDATA-31</td><td>461:460</td><td>Debug Log Pointer</td><td>2</td><td>This shall be set to Cah if the Latency Monitor Debug Telemetry Log is valid.</td></tr>
<tr><td>LMDATA-32</td><td>463:462</td><td>Debug Counter Trigger Source</td><td>2</td><td>When the Debug Counter Trigger Source bit is set to 1b the debug log is valid, and this is the counter that triggered the debug log. When the Debug Counter Trigger Source bit is cleared to 0b this is not the counter that triggered the debug log. No more than 1 bit in this field shall be set. When the Debug Counter Trigger Source is 0b the Debug Log Latency,</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>Debug Log Latency Stamp, Debug Log Pointer, and Debug Counter Trigger Source is not valid.
If the debug log has been cleared this field shall be cleared to zero.
<table>
<tr><th>Bit</th><th>Bucket</th><th>Counter</th><th>Bit Description</th></tr>
<tr><td>15:12</td><td>N/A</td><td>N/A</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>11</td><td>3</td><td>De-allocate/TRIM</td><td>Debug Counter Trigger Source 11</td></tr>
<tr><td>10</td><td>3</td><td>Write</td><td>Debug Counter Trigger Source 10</td></tr>
<tr><td>9</td><td>3</td><td>Read</td><td>Debug Counter Trigger Source 9</td></tr>
<tr><td>8</td><td>2</td><td>De-allocate/TRIM</td><td>Debug Counter Trigger Source 8</td></tr>
<tr><td>7</td><td>2</td><td>Write</td><td>Debug Counter Trigger Source 7</td></tr>
<tr><td>6</td><td>2</td><td>Read</td><td>Debug Counter Trigger Source 6</td></tr>
<tr><td>5</td><td>1</td><td>De-allocate/TRIM</td><td>Debug Counter Trigger Source 5</td></tr>
<tr><td>4</td><td>1</td><td>Write</td><td>Debug Counter Trigger Source 4</td></tr>
<tr><td>3</td><td>1</td><td>Read</td><td>Debug Counter Trigger Source 3</td></tr>
<tr><td>2</td><td>0</td><td>De-allocate/TRIM</td><td>Debug Counter Trigger Source 2</td></tr>
<tr><td>1</td><td>0</td><td>Write</td><td>Debug Counter Trigger Source 1</td></tr>
<tr><td>0</td><td>0</td><td>Read</td><td>Debug Counter Trigger Source 0</td></tr>
</table></td></tr>
<tr><td>LMDATA-33</td><td>464</td><td>Debug Log Stamp Units</td><td>1</td><td><table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:1</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td>0</td><td>When set to 1b the Debug Latency Stamp was based on receiving the NVMe Timestamp and the offset from this. When cleared to 0b the Debug Latency Stamp was based on power on hours since the NVMe Timestamp was not received.</td></tr></table></td></tr>
<tr><td>LMDATA-34</td><td>493:465</td><td>Reserved</td><td>29</td><td>Shall be cleared to zero.</td></tr>
<tr><td>LMDATA-35</td><td>495:494</td><td>Log Page Version</td><td>2</td><td>This indicates the version of this log page. Shall be set to 0004h.</td></tr>
<tr><td>LMDATA-36</td><td>511:496</td><td>Log Page GUID</td><td>16</td><td>Shall be set to 85D45E58D4E643709C6C84D08CC07A92h.</td></tr>
</table>

4.8.10 Device Capabilities (Log Identifier C4h) Requirements

This log provides the host with a consolidated report of critical device-specific support information. This vendor-specific log page, C4h shall be 4096 bytes with the following functional requirements and field format:

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>DCLP-1</td><td>1:0</td><td>PCI Express Ports</td><td>2</td><td>The number of physical PCI Express ports supported by the device.</td></tr>
<tr><td>DCLP-2</td><td>3:2</td><td>OOB Management Support</td><td>2</td><td>Bit field indicating the OOB Management interfaces supported by the device.<table><tr><th>Bit</th><th>Bit Description</th></tr><tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the OOB Management requirements of this specification.</td></tr><tr><td>14:3</td><td>Reserved. Shall be cleared to zero.</td></tr><tr><td>2</td><td>Shall be set to 1b if NVMe Basic Management Command is supported.</td></tr><tr><td>1</td><td>Shall be set to 1b if MCTP over PCIe VDM is supported.</td></tr><tr><td>0</td><td>Shall be set to 1b if MCTP over SMBus is supported.</td></tr></table></td></tr>
<tr><td>DCLP-3</td><td>5:4</td><td>Write Zeroes Command Support</td><td>2</td><td>Bit field indicating the Write Zeroes command requirements supported by the device.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the Write Zeroes command requirements of this specification.</td></tr>
<tr><td>14:5</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>4</td><td>Shall be set to 1b if the requirements of NVMe-IO-6 are met.</td></tr>
<tr><td>3</td><td>Shall be set to 1b if the requirements of NVMe-IO-5 are met.</td></tr>
<tr><td>2</td><td>Shall be set to 1b if setting the FUA bit is supported.</td></tr>
<tr><td>1</td><td>Shall be set to 1b if setting the DEAC bit is supported.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Write Zeroes command is supported.</td></tr>
</table>
</td></tr>
<tr><td>DCLP-4</td><td>7:6</td><td>Sanitize Command Support</td><td>2</td><td>Bit field indicating the Sanitize command requirements supported by the device.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the Sanitize command requirements of this specification.</td></tr>
<tr><td>14:5</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>4</td><td>Shall be set to 1b if Deallocate LBAs is supported.</td></tr>
<tr><td>3</td><td>Shall be set to 1b if Overwrite is supported.</td></tr>
<tr><td>2</td><td>Shall be set to 1b if Block Erase is supported.</td></tr>
<tr><td>1</td><td>Shall be set to 1b if Crypto-Erase is supported.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Sanitize command is supported.</td></tr>
</table>
</td></tr>
<tr><td>DCLP-5</td><td>9:8</td><td>Dataset Management Command Support</td><td>2</td><td>Bit field indicating the Dataset Management command requirements supported by the device.
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td></td><td>Dataset Management command requirements of this specification.</td></tr>
<tr><td>14:2</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>1</td><td>Shall be set to 1b if Attribute – Deallocate (AD) is supported.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Dataset Management command is supported.</td></tr>
</table></td></tr>
<tr><td>DCLP-6</td><td>11:10</td><td>Write Uncorrectable Command Support</td><td>2</td><td>Bit field indicating the Write Uncorrectable command requirements supported by the device.<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the Write Uncorrectable command requirements of this specification.</td></tr>
<tr><td>14:4</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>3</td><td>Shall be set to 1b if the SMART / Health Information (Log Identifier 02h) requirements of NVMe-IO-14 are met.</td></tr>
<tr><td>2</td><td>Shall be set to 1b if the device supports setting uncorrectable on the maximum number of LBAs allowed by the NVMe specification.</td></tr>
<tr><td>1</td><td>Shall be set to 1b if the device supports setting uncorrectable on a single LBA.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Write Uncorrectable command is supported.</td></tr>
</table></td></tr>
<tr><td>DCLP-7</td><td>13:12</td><td>Fused Operation Support</td><td>2</td><td>Bit field indicating the fused command pairs requirements supported by the device.<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>15</td><td>Shall be set to 1b if the device has been tested and found to comply with the fused command pair support requirements of this specification.</td></tr>
<tr><td>14:1</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Compare and Write fused command pair is supported.</td></tr>
</table></td></tr>
<tr><td>DCLP-8</td><td>15:14</td><td>Minimum Valid DSSD Power State</td><td>2</td><td>Shall be set to the lowest numbered valid DSSD Power State. Setting a DSSD Power State less than</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				this value results in Invalid Field in Command (see DSSDPSS-2).
DCLP-9	143:16	DSSD Power State Descriptors	128	Byte 16 is reserved and shall be cleared to zero. The device shall populate bytes 17-143 with a single-byte DSSD Power State Descriptor for each DSSD Power State (see Section 4.8.11.1 DSSD Power State Descriptor) starting with DSSD Power State 1. The DSSD Power State Descriptors shall be populated in increasing order of power state number, i.e., byte 17 contains the descriptor for DSSD Power State 1, byte 18 contains the descriptor for DSSD Power State 2, etc. up to byte 143 which contains the descriptor for DSSD Power State 127.
DCLP-10	4077:144	Reserved	3934	Shall be cleared to zero.
DCLP-11	4079:4078	Log Page Version	2	This indicates the version of the mapping this log page uses. Shall be set to 0001h.
DCLP-12	4095:4080	Log Page GUID	16	Shall be set to B7053C914B58495D98C9E1D10D054297h.

4.8.11 DSSD Power State Requirements

DSSD power states provide an alternative interface to NVMe power management. DSSD and NVMe Power States are numbered in opposite directions – higher numbered DSSD Power States consume more power than lower numbered DSSD Power States, whereas higher numbered NVMe Power States consume less power than lower numbered NVMe Power States.

Selecting a DSSD Power State via a Set Features command (see Section 4.15.15 DSSD Power State (Feature Identifier C7h) Set Feature) causes the device to run at the highest powered NVMe Power State whose Maximum Power (MP) in watts is less than or equal to the number of the DSSD Power State. For example, placing the device in DSSD Power State number 20 causes the device to enter the highest powered NVMe Power State whose Maximum Power (MP) is less than or equal to 20 watts.

In general, a device will have more DSSD Power States than NVMe Power States. Selecting a DSSD Power State whose number is not the Maximum Power (MP) in watts of an NVMe Power State causes the device to drop down to run at the highest powered NVMe Power State whose Maximum Power (MP) does not exceed the number of the DSSD Power State in watts.

For example, suppose that a device supports NVMe Power State 5 that consumes no more than 20 watts Maximum Power (MP) and NVMe Power State 6 that consumes no more than 16 watts Maximum Power (MP). Selecting DSSD Power State 18 (i.e., 18 watts Maximum Power (MP)) causes the device to enter NVMe Power State 6 and consume no more than 16 watts Maximum Power (MP). In this example the device reports 18 as its DSSD Power State and 6 as its NVMe Power State. Selecting NVMe Power State 6 causes the device to report 16 as its DSSD Power State and 6 as its NVMe Power State.

DSSD Power State 0 does not exist, as a device that is consuming 0 watts Maximum Power (MP) is powered off. DSSD Power States that represent a Maximum Power (MP) less than the lowest powered NVMe Power State are invalid. The lowest numbered valid DSSD Power State is indicated in DCLP-8.

4.8.11.1 DSSD Power State Descriptor

The DSSD Power State Descriptor for each DSSD Power State has the following format:

Requirement ID	Bits	Field	Field Description
DSSDPSD-1	7	Valid DSSD Power State	Shall be set to 1b if the number of this DSSD Power State is greater than or equal to the Minimum Valid OCP Power State (see DCLP-8).
DSSDPSD-2	6:5	Reserved	Shall be cleared to zero.
DSSDPSD-3	4:0	NVMe Power State	Shall be set to the number of the highest powered NVMe Power State whose Maximum Power (MP) in watts is less than or equal to the number of this DSSD Power State.

4.8.12 Unsupported Requirements (Log Identifier C5h)

This log provides a host with a consolidated report of all the Requirement IDs that the device does not support. This vendor-specific log page, C5h shall be 4096 bytes with the following functional requirements and field format:

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>URLP-1</td><td>1:0</td><td>Unsupported Count</td><td>2</td><td>The number of Unsupported Requirement IDs.</td></tr>
<tr><td>URLP-2</td><td>15:2</td><td>Reserved</td><td>14</td><td>Shall be cleared to zero.</td></tr>
<tr><td>URLP-3</td><td>4063:16</td><td>Unsupported Requirements List</td><td>4048</td><td>This structure shall be populated with between 0 and 253 16-byte zero padded ASCII strings. Each string shall be the Requirement ID of a requirement that the device does not support. Entries shall be in alphabetical order. Unused entries shall be cleared to zero.
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>31:16</td><td>First Unsupported Requirement ID entry.</td></tr>
<tr><td>47:32</td><td>Second Unsupported Requirement ID entry.</td></tr>
<tr><td>…</td><td>…</td></tr>
<tr><td>4063:4048</td><td>253rd Unsupported Requirement ID entry.</td></tr>
</table></td></tr>
<tr><td>URLP-4</td><td>4077:4064</td><td>Reserved</td><td>14</td><td>Shall be cleared to zero.</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
URLP-5	4079:4078	Log Page Version	2	This indicates the version of the mapping this log page uses. Shall be set to 0001h.
URLP-6	4095:4080	Log Page GUID	16	Shall be set to C7BB98B7D0324863BB2C23990E9C722Fh.

4.8.13 Hardware Component (Log Identifier C6h)

This log provides the host with information regarding the hardware components on the device. This log is intended to enable customers to be able to parse this log when there are component issues to be able to find failing devices. This also enables customers to be able to parse this log for failing devices and not require device serial numbers to find the devices with component issues.

Requirement ID	Description
HWCOMP-1	A value of zero shall only be used if the corresponding component data is not available.
HWCOMP-2	Manufacturer codes shall be clearly documented and disclosed to the customer.
HWCOMP-3	This log shall contain the controller ASIC, PMICs, DRAM, NAND, PCB, and all other critical hardware components.
HWCOMP-4	This log shall contain the Device Serial Number.
HWCOMP-5	This log shall contain the Country of Origin.
HWCOMP-6	This log shall contain the Global Hardware Revision.
HWCOMP-7	The first Component Descriptor shall start at byte 64.
HWCOMP-8	The last Component Descriptor shall end at the byte address which is equal to the Hardware Component Log Size – 1
HWCOMP-9	There shall be no Dwords between two Component Descriptors.

4.8.13.1 Component Descriptor

The Hardware Component Log uses a component descriptor of the following format:

Requirement ID	Byte Address	Field	# of Bytes	Field Description
COMPD-1	7:0	Component Date/Lot Size	8	Component Date/Lot Size in Dwords. This is a ones-based value. A value of 0h means the Component Date/Lot Code field does not exist in the descriptor. A value of 1h means this field is 1 Dword in size.
COMPD-2	15:8	Component Additional Information Size	8	Component Additional Information Size in Dwords. This is a ones-based value. A value of 0h means the Component Additional Information field does not exist in the descriptor. A value of 1h means this field is 1 Dword in size.
COMPD-3	19:16	Component Identifier	4	This is the type of component. This shall follow the Component Identifier Table.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
COMPD-4	27:20	Component Manufacture	8	The value in this field shall represent Component Manufacture. If this field is 0h then there is no component manufacture information provided. This value is in hex.
COMPD-5	35:28	Component Revision	8	The value in this field shall represent the Component Revision. If this field is 0h then there is no Component Revision information provided. This value is in hex.
COMPD-6	43:36	Component Manufacture Code	8	The value in this field shall represent the Component Manufacture Code. If this field is 0h then there is no Component Manufacture Code information provided. This value is in hex.
COMPD-7	(Component Date/Lot Size * 4) + 43:44	Component Date/Lot Code	Component Date/Lot Size * 4	The value in this field shall represent a Component Date/Lot Code. If this field is 0h then there is no Component Date/Lot code provided. This value is in hex.
COMPD-8	(Component Additional Information Size * 4) + (Component Date/Lot Size * 4) + 43:Component Date/Lot Size * 4) + 44	Component Additional Information	Component Additional Information Size * 4	The value In this field shall represent additional component information. If this field is 0h then there is no additional component information. This value is in hex unless otherwise noted in the component identifier.

4.8.13.2 Component Identifier Table

Below is a table of the Component Identifiers.

Requirement ID	Component Identifier	Description
CITYPE-1	0000h	Reserved
CITYPE-2	0001h	This is a Controller ASIC component
CITYPE-3	0002h	This is a NAND Component.
CITYPE-4	0003h	This is a DRAM Component.
CITYPE-5	0004h	This is a PMIC Component.
CITYPE-6	0005	This is a PCB Component
CITYPE-7	0006	This is a capacitor component.

Requirement ID	Component Identifier	Description
CITYPE-8	0007	This is a resistor component.
CITYPE-9	0008	This is a case component.
CITYPE-10	0009	Device Serial Number. The serial number shall go into the Additional Component Information field and the other fields shall be 0h, except for the Component Identifier and Component Additional Information Size.
CITYPE-11	000A	Country of Origin. This is a UTF-8 string detailing the device Country of Origin. The Country of Origin value shall go into the Additional Component Information field and the other fields shall be 0h, except for the Component Identifier and Component Additional Information Size.
CITYPE-12	000B	Global Device Hardware Revision. The Global Device Hardware Revision shall go into the Additional Component Information field and the other fields shall be 0h, except for the Component Identifier and Component Additional Information Size.
CITYPE-13	7FFF-000C	Reserved
CITYPE-14	8000-FFFF	Vendor Unique Component

4.8.13.3 Hardware Component Log

Requirement ID	Byte Address	Field	# of Bytes	Field Description
HWCLOG-1	1:0	Log Page Version	2	This indicates the version of the mapping this log page uses. Shall be set to 0001h.
HWCLOG-2	15:2	Reserved	14	Shall be cleared to zero.
HWCLOG-3	31:16	Log Page GUID	16	Shall be set to BCB6821F30CD4ED0B76B31B99F0F57DC.
HWCLOG-4	47:32	Hardware Component Log Size	16	This is the size of the Hardware Component Log in Dwords.
HWCLOG-5	63:48	Reserved	16	Reserved. Shall be cleared to 0h.
HWCLOG-6	Hardware Component Log Size – 1:64	Component Descriptions	Hardware Component Log Size – 64	This shall contain the hardware descriptors.

4.8.14 TCG Configuration (Log Identifier C7h)

TCG Configuration Log shall be 512 bytes in size. This log defines the data points related to TCG Opal Security Subsystem Class (SSC) that would be useful to collect for diagnostic purposes. These data points give point-in-time insight into the configuration and runtime state of the device which can be used to explain either a device or host behavior.

TCG Log Persistence Type	Persistence Type
A	Persistent across TCG revert: PSID, Admin SP, Locking SP
B	Not persistent across TCG revert

Requirement ID	Byte Address	Field	# of Bytes	Field Description	Persistence Type	TCG Log Persistence Type
TCGL-1	0	State	1	Bit 0: C_PIN_SID Value indicates whether C_PIN_SID object's PIN column value is equal to MSID or not. 0 = MSID value 1 = Not MSID value Bit 1: C_PIN_SID Authentication Blocked indicates whether authentication of C_PIN_SID authority is blocked or not due to Block SID Authentication command. 0 = Not blocked 1 = Blocked Bit 2: Locking Enabled (feature code 0x0002) indicates whether the Locking SP is currently in 'Manufactured-Inactive' or 'Manufactured' state. 0 = Locking SP not activated i.e., Manufactured-Inactive state 1 = Locking SP activated i.e., Manufactured state Bit 3: Single User Mode Owner **i**ndicates the ownership policy for the RangeStart and RangeLength columns of all Locking Objects in Single User Mode. 0 = User Authority 1 = Admins Authority All other bits shall be cleared to zero.	1	N/A
TCGL-2	3:1	Reserved	3	Shall be cleared to zero.	N/A	N/A

Requirement ID	Byte Address	Field	# of Bytes	Field Description	Persistence Type	TCG Log Persistence Type
TCGL-3	4	Locking SP Activation Count	1	Count of transitions of the Locking SP from ‘Manufactured-Inactive’ to ‘Manufactured’ since the last power cycle.	2	A
TCGL-4	5	Tper Revert Count	1	Count of successful invocations of the Revert method on the Admin SP which causes the entire Tper to revert to Original Factory State, since the last power cycle.	2	A
TCGL-5	6	Locking SP Revert Count	1	Count of successful invocations of the Revert method on the Locking SP which causes the Locking SP to transition from ‘Manufactured’ to ‘Manufactured-Inactive’, since the last power cycle.	2	A
TCGL-6	7	Number of Locking Objects	1	The number of Locking Objects supported in the Locking table including the Global Locking Object. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-7	8	Number of Single User Mode Locking Objects	1	The current number of Locking Objects in the Locking table that are in Single User Mode including the Global Locking Object. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-8	9	Number of Range Provisioned Locking Objects	1	The current number of non-Global Locking Objects with non-zero RangeLength value. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-9	10	Number of Namespace Provisioned Locking Objects	1	The current number of non-Global Locking Objects with non-zero NamespaceID value. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-10	11	Number of Read Locked Locking Objects	1	The current number of provisioned non-Global Locking Objects with ReadLockEnabled set to True and ReadLocked set to True. If	1	N/A

Requirement ID	Byte Address	Field	# of Bytes	Field Description	Persistence Type	TCG Log Persistence Type
				the Locking SP is inactive, then this value shall be cleared to zero.		
TCGL-11	12	Number of Write Locked Locking Objects	1	The current number of provisioned non-Global Locking Objects with WriteLockEnabled set to True and WriteLocked set to True. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-12	13	Number of Read Unlocked Locking Objects	1	The current number of provisioned non-Global Locking Objects with ReadLockEnabled set to True and ReadLocked set to False. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-13	14	Number of Write Unlocked Locking Objects	1	The current number of provisioned non-Global Locking Objects with WriteLockEnabled set to True and WriteLocked set to False. If the Locking SP is inactive, then this value shall be cleared to zero.	1	N/A
TCGL-14	15	Reserved	1	Shall be cleared to zero.	N/A	N/A
TCGL-15	19:16	SID Authentication Try Count	4	The Tries value for the C_PIN_SID Credential Object which indicates the number of failed authentication attempts for that object since the last power cycle, successful authentication or Tper revert.	1	N/A
TCGL-16	23:20	SID Authentication Try Limit	4	The TryLimit value for the C_PIN_SID Credential Object which indicates the maximum number of failed authentication attempts for that object.	1	N/A
TCGL-17	27:24	Programmatic TCG Reset Count	4	The count of Programmatic TCG Resets received by the device since the last power cycle.	3	A
TCGL-18	31:28	Programmatic Reset Lock Count	4	The count of state transitions for all provisioned non-Global Locking Objects in the following conditions: • Read Lock Transition when ReadLockedEnable = True and	3	A

Requirement ID	Byte Address	Field	# of Bytes	Field Description	Persistence Type	TCG Log Persistence Type
				ReadLocked transitions from False to True. - Write Lock Transition when WriteLockedEnable = True and WriteLocked transitions from False True. The count shall increment for both explicit and implicit Programmatic Resets since the last power cycle.		
TCGL-19	35:32	TCG Error Count	4	The count of TCG errors reported to the host.	3	A
TCGL-19	493:36	Reserved	458	Shall be cleared to zero.	N/A	N/A
TCGL-20	495:494	Log Version	2	Shall be set to 0001h.	N/A	N/A
TCGL-21	511:496	Log Page GUID	16	Shall be set to 54E02A9DFA5447C083E6E07EBD244006.	N/A	N/A

4.8.15 Telemetry String Log (Log Identifier C9h)

The Telemetry String Log enables mapping of Statistic Identifiers, Debug Class/Event Identifiers and Debug Class/VU Event Identifier combinations to ASCII strings.

4.8.15.1 Telemetry String Log Requirements

Requirement ID	Description
TELSLG-1	All Statistic Identifiers contained in Data Area 1 and Data Area 2 shall be contained in this log.
TELSLG-2	String Identifier Table entries shall be in numeric order such that the smallest Vendor Unique Statistic Identifier is at the lowest address.
TELSLG-3	All Debug Event Class/ Event Identifier combinations contained in Telemetry Data Area 1 and Telemetry Data Area 2 shall be contained in this log.
TELSLG-4	All Debug Event Class/ VU Event Identifier combinations contained in Telemetry Data Area 1 and Telemetry Data Area 2 shall be contained in this log.
TELSLG-5	String Event Table entries shall be in numeric order such that the smallest Debug event Class is at the lowest address.
TELSLG-6	String VU Event Table entries shall be in numeric order such that the smallest Debug event Class is at the lowest address.
TELSLG-7	String Event Table entries shall be in numeric order such that for a single Debug/Event Identifier combination the smallest Event Identifier is at the lowest address.
TELSLG-8	String VU Event Table entries shall be in numeric order such that for a single Debug/VU Event Identifier combination the smallest VU Event Identifier is at the lowest address.

Requirement ID	Description
TELSLG-9	In the ASCII table characters which are not valid (beyond the length of the ASCII length) shall be set to space <space>.
TELSLG-10	The String Identifier Table shall have no gaps between this table and the String Event Table.
TELSLG-11	The String Event Table entry shall have no gaps between this table and the VU String Event Table.
TELSLG-12	The VU String Event Table and the ASCII table shall have no gaps between these tables.
TELSLG-13	The Sting Identifier Table Entries, The String Event Table Entries and the String VU Event Table Entries may point to the same ASCII Table Entry.

4.8.15.2 Telemetry String Log Format

The Telemetry Sting log enables mapping data constructs in Telemetry Data Area 1 and Telemetry Data Area 2 to ASCII strings. The size of the Telemetry String Log can be found in Telemetry Data Area 1 or in the Telemetry String log. The following is the format of the Telemetry String Log:

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELSTR-LOG-1	0	Log Page Version	1	Shall be set to 01h.
TELSTR-LOG-2	15:1	Reserved	15	Shall be cleared to zero.
TELSTR-LOG-3	31:16	Log Page GUID	16	Shall be set to B13A83691A8F408B9EA495940057AA44h.
TELSTR-LOG-4	39:32	Telemetry String Log Size (SLS)	8	Shall be set to the number of Dwords in the String Log.
TELSTR-LOG-5	63:40	Reserved	24	Shall be cleared to zero.
TELSTR-LOG-6	71:64	Statistics Identifier String Table Start (SITS)	8	Statistics Identified String Table Start (SITS) shall be set to the number of Dwords from byte 0 of this log page to the start of the Statistics Identifier String Table. This shall be set to 6Ch.
TELSTR-LOG-7	79:72	Statistics Identifier String Table Size (SITSZ)	8	Statistics Identifier String Table Size (SITSZ) shall be set to the number of Dwords in the Statistics Identifier String Table. If there are no vendor unique defined values this shall be set to 0h.
TELSTR-LOG-8	87:80	Event String Table Start (ESTS)	8	Event String Table Start (ESTS) Shall be set to the number of Dwords from byte 0 of this log page to the start of the Event String Table.
TELSTR-LOG-9	95:88	Event String Table Size (ESTSZ)	8	Event String Table Size (ESTSZ) shall be set to the number of Dwords in the Event String Table.
TELSTR-LOG-10	103:96	VU Event String Table Start	8	VU Event String Table Start. Shall be set to the number of Dwords from byte 0 of this log page to the start of the VU Event String Table.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELSTR-LOG-11	111:104	VU Event String Table Size	8	VU Event String Table Size (VUESTZ) shall be set to the number of Dwords in the VU Event String Table.
TELSTR-LOG-12	119:112	ASCII Table Start (ASCTS)	8	ASCII Table Start (ASCTS). This is the number of Dwords from byte 0 of this log page until the ASCII Table Starts.
TELSTR-LOG-13	127:120	ASCII Table Size (ASCTSZ)	8	ASCII Table Size (ASCTSZ). This is the number of Dwords in the ASCII Table.
TELSTR-LOG-14	143:128	FIFO 1 ASCII String	16	FIFO 0 ASCII String. This is the name of FIFO 0. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-15	159:144	FIFO 2 ASCII String	16	FIFO 1 ASCII String. This is the name of FIFO 1. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-16	175:160	FIFO 3 ASCII String	16	FIFO 2 ASCII String. This is the name of FIFO 2. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-17	191:176	FIFO 4 ASCII String	16	FIFO 3 ASCII String. This is the name of FIFO 3. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-18	207:192	FIFO 5 ASCII String	16	FIFO 4 ASCII String. This is the name of FIFO 4. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-19	223:208	FIFO 6 ASCII String	16	FIFO 5 ASCII String. This is the name of FIFO 5. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-20	239:224	FIFO 7 ASCII String	16	FIFO 6 ASCII String. This is the name of FIFO 6. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-21	255:240	FIFO 8 ASCII String	16	FIFO 7 ASCII String. This is the name of FIFO 7. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-22	271:256	FIFO 9 ASCII String	16	FIFO 8 ASCII String. This is the name of FIFO 8. If the FFO has no string name this shall be populated with 0h.
TELSTR-LOG-23	287:272	FIFO 10 ASCII String	16	FIFO 9 ASCII String. This is the name of FIFO 9. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-24	303:288	FIFO 11 ASCII String	16	FIFO 10 ASCII String. This is the name of FIFO 10. If the FIFO has no string name this shall be populated with 0h.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELSTR-LOG-25	319:304	FIFO 12 ASCII String	16	FIFO 11 ASCII String. This is the name of FIFO 11. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-26	335:320	FIFO 13 ASCII String	16	FIFO 12 ASCII String. This is the name of FIFO 12. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-27	351:336	FIFO 14 ASCII String	16	FIFO 13 ASCII String. This is the name of FIFO 13. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-28	367:352	FIFO 15 ASCII String	16	FIFO 14 ASCII String. This is the name of FIFO 14. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-29	383:368	FIFO 16 ASCII String	16	FIFO 15 ASCII String. This is the name of FIFO 15. If the FIFO has no string name this shall be populated with 0h.
TELSTR-LOG-30	431:384	Reserved	48	Reserved. Shall be cleared to 0h.
TELSTR-LOG-31	(SITSZ *4) + 431:432	Statistics Identifier String Table	SITSZ * 4	The format of the Statistics Identifier String Table can be seen below.
TELSTR-LOG-32	(ESTSZ * 4) + (SITSZ * 4) + 431:SITSZ * 4) + 432	Event String Table	ESTSZ * 4	The format of the Event String Table can be seen below.
TELSTR-LOG-33	(VUEDH * 4) + (ESTSZ * 4) + (SITSZ * 4) + 431:ESTSZ * 4) + (SITSZ * 4) + 432	VU Event String Table	VUESTZ * 4	The format of the VU Event String Table can be seen below.
TELSTR-LOG-34	(ASCTSZ * 4) + (VUEDH * 4) + (ESTSZ * 4) + (SITSZ * 4) + 432: (VUEDH *	ASCII Table	ASCTSZ * 4	The format of the ASCII can be seen below.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
	4) + (ESTSZ * 4) + (SITSZ * 4) + 432			

A diagram of the String Log format can be seen below.

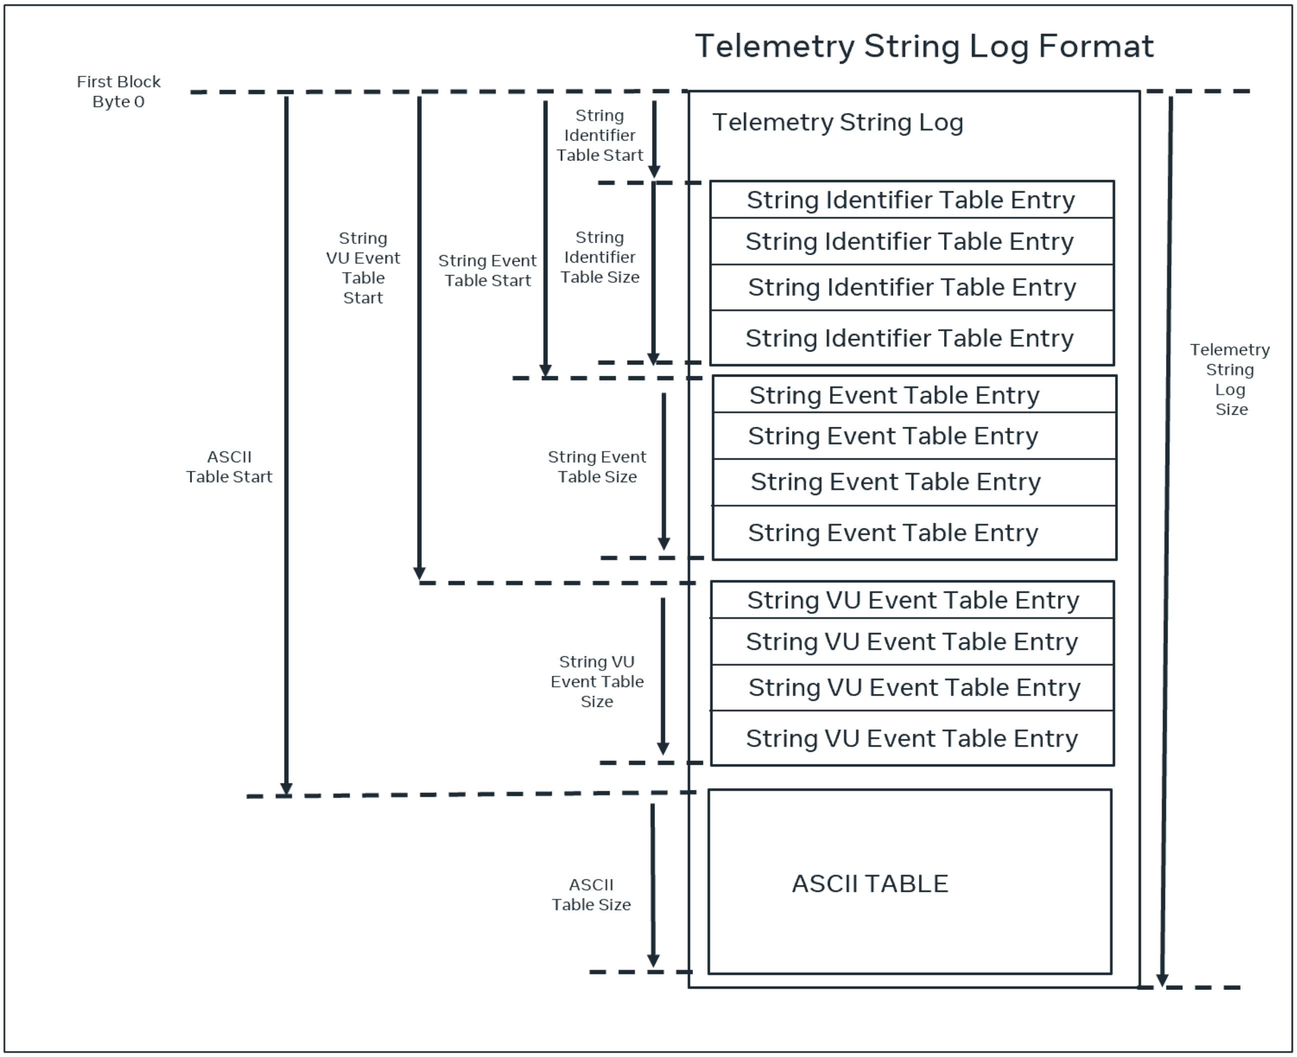
Telemetry String Log Format
First Block Byte 0
String VU Event Table Start
String Event Table Start
String Identifier Table Start
String Identifier Table Size
ASCII Table Start
String Event Table Size
String VU Event Table Size
ASCII Table Size
Telemetry String Log
String Identifier Table Entry
String Identifier Table Entry
String Identifier Table Entry
String Identifier Table Entry
String Event Table Entry
String Event Table Entry
String Event Table Entry
String Event Table Entry
String VU Event Table Entry
String VU Event Table Entry
String VU Event Table Entry
String VU Event Table Entry
ASCII TABLE
Telemetry String Log Size

4.8.15.3 Statistics Identifier String Table Entry

The Statistic Identifier String Table enables translating from vendor unique statistic identifiers to ASCII labels. The format of an entry for the Statistic Identifier String Table Entry is below.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
STR-ID-1	1:0	Vendor Specific Statistic Identifier	2	Shall be set the Vendor Unique Statistic Identifier number.
STR-ID-2	2	Reserved	1	Shall be cleared to zero.
STR-ID-3	3	ASCII ID Length	1	Shall be set the number of ASCII Characters that are valid. This is a zero-based value so a value of 0h is 1 character.
STR-ID-4	11:4	ASCII ID offset	8	Shall be set to the offset from Dword 0/Byte 0 of the Start of the ASCII Table to the first character of the string for this Statistic Identifier string. This value is in Dwords.
STR-ID-5	15:12	Reserved	4	Shall be cleared to zero.

4.8.15.4 Event Identifier String Table Entry

The Event Identifier String Table enables translating from Event Identifier to ASCII labels. The format of an entry for the Event Identifier String Table Entry is below.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
STR-EV-1	0	Debug Event Class	1	Shall be set the Debug Class.
STR-EV-2	2:1	Event Identifier	2	Shall be set to the Event Identifier.
STR-EV-3	3	ASCII ID Length	1	Shall be set to the number of ASCII Characters that are valid. This is a zero-based value so a value of 0h is 1 character.
STR-EV-4	11:4	ASCII ID offset	8	This is the offset from Dword 0/ byte 0 of the start of the ASCII table to the ASCII data for this identifier. This value is in Dwords.
STR-EV-5	15:12	Reserved	4	Shall be cleared to zero.

4.8.15.5 VU Event Identifier String Table Entry

The VU Event Identifier String Table enables translating from VU Event Identifier to ASCII labels. The format of an entry for the VU Event Identifier String Table Entry is below.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
STR-VUEV-1	0	Debug Event Class	1	Shall be set the Debug Class.
STR-VUEV-2	2:1	VU Event Identifier	2	Shall be set to the VU Event Identifier

Requirement ID	Byte Address	Field	# of Bytes	Field Description
STR-VUEV-3	3	ASCII ID Length	1	Shall be set to the number of ASCII Characters that are valid. This is a zero-based value so a value of 0h is 1 character.
STR-VUEV-4	11:4	ASCII ID offset	8	This is the offset from Dword 0/ byte 0 of the start of the ASCII table to the ASCII data for this identifier. This value is in Dwords.
STR-VUEV-5	15:12	Reserved	4	Shall be cleared to zero.

4.8.15.6 ASCII String Example

Below is an example of how the ASCII ID Length/Decode example: ASCII ID Length = 0Ch

Byte	Data Byte	ASCII Character
400	54h	T
401	45h	E
402	53h	S
403	54h	T
404	20h	<space>
405	45h	E
406	58h	X
407	41h	A
408	4Dh	M
409	50h	P
410	4Ch	L
411	45h	E
412	21h	!
413	20h	space
414	20h	space
415	20h	space

In this example the length is 0Ch, which is 12 decimal. This is a zero-based count and thus there are 13 characters to print. The first 13 characters shall be printed. Thus, the message to print is:

"TEST EXAMPLE!"

The ASCII characters of space are not printed since the offsets are Dword, these strings are padding and are not part of the next string.

4.9 Host and Controller Initiated Telemetry Logs

The following requirements define a standardized format for Telemetry Data Area 1 and Telemetry Data Area 2 for both Telemetry Host-Initiated Log (Log Identifier 07h) and Telemetry Controller-Initiated Log (Log Identifier 08h).

4.9.1 Theory of Operation

This feature is designed to enable both customers and suppliers to have increased visibility into monitoring and debugging devices.

4.9.2 Data Partitioning Overview

The data provided by the device for debugging is broken into two categories. There is data that impacts I/O and data that does not impact I/O. The data generated by the device for monitoring and/or debugging which does not impact I/O (performance/latency) for active commands in progress goes into Data Area 1. The data generated by the device for monitoring and/or debugging which impacts I/O (performance/latency) for active commands in progress goes into Data Area 2. It should be noted that the data format in the host-initiated telemetry and the controller-initiated log is the same format and the information reported in both host-initiated and controller-initiated Data Areas 1 and 2 shall be the same reported data.

4.9.3 Statistic Area

The Statistic Area contains a group of Statistic Descriptors. The Statistic Descriptors may be counters or other statistics where each statistic is a value with a statistic size. The attributes associated with each of these statistics are Statistic Identifier, Behavior Type and the length of the data field associated with the Statistic Identifier.

4.9.4 Event FIFO

The Event FIFO is a queue that contains events that are generated. There are a total of up to 16 independent Event FIFO queues that may be configured between Telemetry Data Area 1 and Telemetry Data Area 2. It should be noted that the ordering relationship of events in the event FIFO is determined by the device.

4.9.5 Existing NVMe Logs

There are many existing NVMe Logs that are useful for monitoring the device. These logs have been overlayed into Telemetry Data Area 1. This is done to enable all the monitoring information to be gathered when polling Telemetry Data Area 1.

4.9.6 Debug Structure Overview Summary

The debug structure for Data Area 1 and 2 can be seen in the figure below.

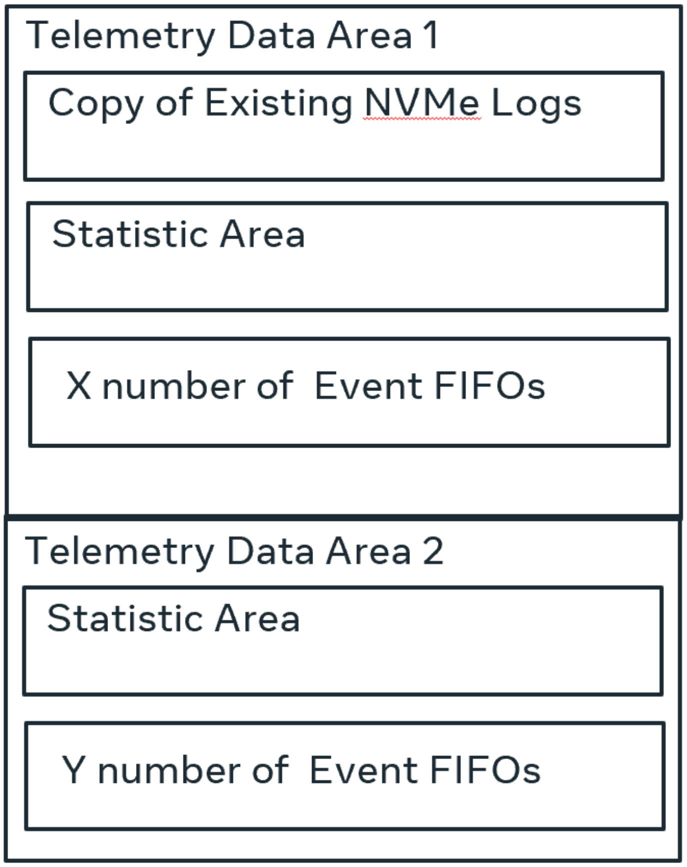

4.9.7 Telemetry Data Requirements

The following applies to telemetry logging as the ability to quickly debug failures is required:

Requirement ID	Description
TEL-1	The device shall track the operational/event history and any critical parameters that can be used to debug issues.
TEL-2	The supplier shall provide a table that categorizes the reason identifiers that are a super set of the panic IDs in EREC-4.
TEL-3	Obsolete. See STD-LOG-23.
TEL-4	The Reason Identifier field (in the Telemetry Controller-Initiated log page and the Telemetry Host-Initiated log page) shall be the most recent failure identifier and shall not be cleared by a power cycle or reset.
TEL-5	The table below provides the specifications for the controller-initiated and the host-initiated log page Data Areas 1 and 2. Any vendor-specific data required for root-cause analysis shall be included in data area 3. If the vendor-specific data required for root-cause analysis does not fit within data area 3, then the additional vendor-specific data shall be included in data area 4.

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Telemetry Data Area</th><th>Purpose</th><th>Profile 0 Data Area Size</th><th>Latency Impact to IO</th></tr>
<tr><td>1</td><td>Periodic logging for monitoring trends/problems</td><td>16384 bytes</td><td>&lt; 1ms typical,
<5 ms max</td></tr>
<tr><td>2</td><td>Logging for debug items that impact latency</td><td>-</td><td>Greater Telemetry Data Area 1 Requirement</td></tr>
<tr><td>3</td><td>Optional Vendor Unique</td><td>-</td><td>-</td></tr>
<tr><td>4</td><td>Optional Vendor Unique</td><td>-</td><td>-</td></tr>
</table>
</td></tr>
<tr><td>TEL-6</td><td>The device shall set the Telemetry Host-Initiated Data Area 1 Last Block field of the Telemetry Host-Initiated Log (Log Identifier 07h) to 0020h as the factory default.</td></tr>
<tr><td>TEL-7</td><td>The device shall set the Telemetry Controller-Initiated Data Area 1 Last Block field of the Telemetry Controller-Initiated Log (Log Identifier 08h) to 0020h as the factory default.</td></tr>
<tr><td>TEL-8</td><td>Information in Telemetry Data Area 1 shall be updated in the background at least once every ten minutes.</td></tr>
<tr><td>TEL-9</td><td>If the device supports an encrypted log for debug this may be contained in Telemetry Data Area 3 and/or 4.</td></tr>
<tr><td>TEL-10</td><td>The number of Timestamp Events inserted into the Event FIFOs shall be sufficient to debug.</td></tr>
<tr><td>TEL-11</td><td>For any issue that causes the controller to generate the Telemetry Controller-Initiated log page, the complete contents of each of the Telemetry Controller-Initiated log page and Telemetry Host-Initiated log page shall contain all of the debug information needed to root cause the issue.</td></tr>
<tr><td>TEL-12</td><td>Updating the Error ID in the TELRI-1 Reason Identifier field shall cause a Controller Initiated Telemetry AEN event.</td></tr>
<tr><td>TEL-13</td><td>When the device encounters a failure/error the Error ID in TELRI-1 shall be a non-zero value.</td></tr>
<tr><td>TEL-14</td><td>Decode of Telemetry Data in Data Area 1 and Data Area 2 as defined in this specification shall be supported through the OCP based NVMe CLI tool. And shall meet the requirements for SEC-22.</td></tr>
<tr><td>TEL-15</td><td>Telemetry Data Area 1 and Telemetry Data Area 2 shall be supported per this specification.</td></tr>
<tr><td>TEL-16</td><td>A minimum of two Telemetry Profiles shall be supported.</td></tr>
<tr><td>TEL-17</td><td>Telemetry Profile 0 shall be the default profile from the factory.</td></tr>
<tr><td>TEL-18</td><td>Telemetry Profile 1 shall be sized for maximum debuggability.</td></tr>
<tr><td>TEL-19</td><td>The device shall support TP4109a to enable the host to access specific host telemetry.</td></tr>
<tr><td>TEL-20</td><td>On a firmware update the Vendor Specific Statistics Identifier, the Debug Event Class/Events Identifier combinations and the Debug Class/VU Event Identifier combinations shall not be re-purposed or removed from the telemetry string log. This is to ensure on a firmware update that any existing telemetry log is able to be decoded correctly.</td></tr>
<tr><td>TEL-21</td><td>Telemetry Controller-Initiated log pages shall only be generated (Telemetry Controller-Initiated Data Available transitions from 00h to 01h) for error conditions (e.g., they shall not be generated periodically, or in response to Percentage Used changes, etc.).</td></tr>
</table>

Requirement ID	Description
TEL-22	The Telemetry Controller-Initiated log page and Telemetry Host-Initiated log page shall be NVM subsystem in scope.

4.9.8 Telemetry Reason Identifier Format

The format of the Reason Identifier shall be as follows:

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th># of Bytes</th><th>Field</th><th>Field Description</th></tr>
<tr><td>TELRI-1</td><td>63:0</td><td>64</td><td>Error ID</td><td>The device shall set a unique Error ID defining the persistent or transient error at the time of log capture.</td></tr>
<tr><td>TELRI-2</td><td>71:64</td><td>8</td><td>File ID</td><td>File name or hash of the source file where the error occurred. Clear to zero if not applicable.</td></tr>
<tr><td>TELRI-3</td><td>73:72</td><td>2</td><td>Line Number</td><td>Line Number in the source file where the error occurred. Clear to zero if not applicable.</td></tr>
<tr><td>TELRI-4</td><td>74</td><td>1</td><td>Valid Flags</td><td>
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:4</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>3</td><td>VU Reason Extension Field is valid.</td></tr>
<tr><td>2</td><td>Shall be set to 1b if the Error ID is valid. Shall be cleared to 0b if Error ID is not valid.</td></tr>
<tr><td>1</td><td>Shall be set to 1b if the File ID is valid. Shall be cleared to 0b if File ID is not valid.</td></tr>
<tr><td>0</td><td>Shall be set to 1b if the Line Number is valid. Shall be cleared to 0b if Line Number is not valid.</td></tr>
</table>
</td></tr>
<tr><td>TELRI-5</td><td>95:75</td><td>21</td><td>Reserved</td><td>Shall be cleared to zero.</td></tr>
<tr><td>TELRI-6</td><td>127:96</td><td>32</td><td>VU Reason Extension</td><td>Vendor Unique reason code.</td></tr>
</table>

4.9.9 Debug Profiles

There are many ways to configure Data Area 1, Data Area 2, the Statistics and Event FIFOs. Depending on the problem being investigated different configurations of these fields may be desired. Debug profiles are the mechanism that can address this challenge. The device will list the number of debug profiles supported. The host may then configure the profile based on the problem being investigated.

4.9.10 Statistic Descriptor

The format of the Statistic Descriptor is below.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
STATF-1	1:0	Statistic Identifier	2	Shall contain the identifier that defines the data format of this entry.

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>STATF-2</td><td>2</td><td>Statistic Info</td><td>1</td><td><table><tr><th>Bit</th><th>Bit Description</th></tr><tr><td>7:4</td><td>Reserved. Shall be cleared to zero.</td></tr><tr><td>3:0</td><td>Behavior Type</td></tr></table></td></tr>
<tr><td>STATF-3</td><td>3</td><td>NS INFO</td><td>1</td><td><table><tr><th>Bit</th><th>Bit Description</th></tr><tr><td>7</td><td>Namespace Information Valid. When set to 1b the Namespace Identifier is valid. When cleared to 0b the Namespace Identifier is not valid</td></tr><tr><td>6:0</td><td>Namespace Identifier. When Namespace Information valid is set to 1b this is the Namespace Identifier (NS ID). When Namespace Information is cleared to 0b this field is not valid and shall be cleared to 0h.</td></tr></table></td></tr>
<tr><td>STATF-4</td><td>5:4</td><td>Statistic Data Size</td><td>2</td><td>Shall contain the number of Dwords that is used by the Statistic Specific Data field. Valid values for this field are 0h to FFFFh.</td></tr>
<tr><td>STATF-5</td><td>7:6</td><td>Reserved</td><td>2</td><td>Reserved. Shall be cleared to 0h.</td></tr>
<tr><td>STATF-6</td><td>(Statistic Data Size *4)+7:8</td><td>Statistic Specific Data</td><td>Statistic Data Size *4</td><td>Shall contain data specific extended information for this Statistic Data Entry type if this exists.</td></tr>
</table>

4.9.11 Statistic Identifiers

The definitions of the Statistic Identifiers are below. Statistic Identifiers that provide value to the vendor in debugging the device should be implemented.

Requirement ID	Statistic Identifier	Recommended Behavior Type	# of Dwords	Statistic Specific Data Field	Description
STATI-1	0000h	-	-	-	Reserved
STATI-2	0001h	1	1	Outstanding Admin Commands	Number of outstanding Admin commands submitted by the host which have not been processed by the controller. Pulled from Submission Queue but not yet pushed to the Completion Queue. This count shall not include Asynchronous Event Request (AER) commands.

Requirement ID	Statistic Identifier	Recommended Behavior Type	# of Dwords	Statistic Specific Data Field	Description
STATI-3	0002h	1	1	Host Write Bandwidth	The percentage of the current write bandwidth allowed to the NAND backend due to host writes.
STATI-4	0003h	1	1	GC Write Bandwidth	The percentage of the current write bandwidth allowed to the NAND backend due to internal writes. STATI-3 and STATI-4 shall add to 100%.
STATI-5	0004h	1	1	Active Namespaces	Shall be the number of Active Namespaces attached to the primary controller.
STATI-6	0005h	5	2	Internal Write Workload	Number of outstanding LBAs corresponding to the Internal Write Queue Depth (STATI-8).
STATI-7	0006h	5	2	Internal Read Workload	Number of outstanding LBAs corresponding to the Internal Read Queue Depth (STATI-9).
STATI-8	0007h	5	1	Internal Write Queue Depth	Number of outstanding Write commands submitted by the host which have not been processed by the controller. Pulled from Submission Queue but not yet pushed to the Completion Queue.
STATI-9	0008h	5	1	Internal Read Queue Depth	Number of outstanding Read commands submitted by the host which have not been processed by the controller. Pulled from Submission Queue but not yet pushed to the Completion Queue.
STATI-10	0009h	5	2	Pending Trim LBA Count	From a device perspective this is the current number of LBAs pending completion of the background trim process due to host Dataset Management – Deallocate commands.
STATI-11	000Ah	4	2	Host Trim LBA Request Count	Number of LBAs that have been requested by Dataset Management – Deallocate commands since last Telemetry Log collection.

Requirement ID	Statistic Identifier	Recommended Behavior Type	# of Dwords	Statistic Specific Data Field	Description
STATI-12	000Bh	1	1	Current NVMe Power State	Currently set NVMe Power State Descriptor at the time of this Telemetry Log collection.
STATI-13	000Ch	1	1	Current DSSD Power State	Currently set DSSD Power State Descriptor at the time of this Telemetry Log collection.
STATI-14	000Dh	4	2	Program Fail Count	The number of program operation failure events for the life of the device.
STATI-15	000Eh	4	2	Erase Fail Count	The number of erase operation failure events for the life of the device.
STATI-16	000Fh	4	4	Read Disturb Writes	Number of bytes written due to read disturb relocations for the life of the device.
STATI-17	0010h	4	4	Retention Writes	Number of bytes written due to retention relocation for the life of the device.
STATI-18	0011h	4	4	Wear Leveling Writes	Number of bytes written due to wear leveling for the life of the device.
STATI-19	0012h	4	2	Read Recovery Writes	Number of bytes written due to read recovery for the life of the device.
STATI-20	0013h	1	2	GC Writes	Number of bytes written due to garbage collection since last Telemetry Log collection using by reading either host-initiated or controller-initiated log.
STATI-21	0014h	4	1	SRAM Correctable Count	Total number of correctable errors due to device SRAM single error correction over the device lifetime.
STATI-22	0015h	4	1	DRAM Correctable Count	Total number of correctable errors due to device DRAM single error correction over the life of the device.
STATI-23	0016h	4	1	SRAM Uncorrectable Count	Total number of uncorrectable errors due to device SRAM double error detection.
STATI-24	0017h	4	1	DRAM Uncorrectable Count	Total number of uncorrectable errors due to device DRAM double error detection.

<table>
<thead>
<tr><th>Requirement ID</th><th>Statistic Identifier</th><th>Recommended Behavior Type</th><th># of Dwords</th><th>Statistic Specific Data Field</th><th>Description</th></tr>
</thead>
<tbody>
<tr><td>STATI-25</td><td>0018h</td><td>4</td><td>1</td><td>Data Integrity Error Count</td><td>Total number of data integrity errors due to FTL metadata integrity checks.</td></tr>
<tr><td>STATI-26</td><td>0019h</td><td>4</td><td>1</td><td>Read Retry Error Count</td><td>The number of reads for the device lifetime performed on the flash because of error correction (e.g., Read retries, LDPC iterations, etc.).</td></tr>
<tr><td>STATI-27</td><td>001Ah</td><td>4</td><td>1</td><td>PERST# Events Count</td><td>Number of PERST# events processed by the NVM Subsystem for the lifetime of the device. This count shall only increment if the CC.EN bit is set to 1b.</td></tr>
<tr><td>STATI-28</td><td>001Bh</td><td>4</td><td>2</td><td>Max Die Bad Block</td><td>This information is based on a single die which has the largest number of bad blocks.
<table>
<thead><tr><th>Byte Address</th><th>Byte Description</th></tr></thead>
<tbody>
<tr><td>0</td><td>Worst Die % of bad blocks, where the range is from 0-100%.</td></tr>
<tr><td>1</td><td>Reserved</td></tr>
<tr><td>3:2</td><td>Worst Die Raw Number of bad blocks</td></tr>
<tr><td>7:4</td><td>Reserved. Cleared to 0h.</td></tr>
</tbody>
</table></td></tr>
<tr><td>STATI-29</td><td>001Ch</td><td>4</td><td>2</td><td>Max NAND Channel Bad Block</td><td>This information is based on the dies in a single NAND channel which has the largest number of bad blocks.
<table>
<thead><tr><th>Byte Address</th><th>Byte Description</th></tr></thead>
<tbody>
<tr><td>0</td><td>Worst NAND Channel % of bad blocks, where the range is 0-100%</td></tr>
<tr><td>1</td><td>Reserved</td></tr>
<tr><td>3:2</td><td>Worst NAND Channel Number of bad blocks</td></tr>
</tbody>
</table></td></tr>
</tbody>
</table>

<table>
<tr><th>Requirement ID</th><th>Statistic Identifier</th><th>Recommended Behavior Type</th><th># of Dwords</th><th>Statistic Specific Data Field</th><th>Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td></td><td><table><tr><td>7:4</td><td>Reserved. Cleared to 0h.</td></tr></table></td></tr>
<tr><td>STATI-30</td><td>001Dh</td><td>4</td><td>2</td><td>Minimum NAND Channel Bad Block</td><td>This information is based on the dies in a single NAND channel which has the smallest number of bad blocks.<table><tr><th>Byte Address</th><th>Byte Description</th></tr><tr><td>0</td><td>Best NAND Channel % of bad blocks where the range is 0-100%</td></tr><tr><td>1</td><td>Reserved</td></tr><tr><td>3:2</td><td>Best NAND Channel Number of bad blocks</td></tr><tr><td>7:4</td><td>Reserved. Cleared to 0h.</td></tr></table></td></tr>
<tr><td>STATI-31</td><td>7FFFh:001Eh</td><td>…</td><td>…</td><td>Reserved</td><td>Statistic Identifiers between 001Eh and 7FFFh are reserved for future expansion.</td></tr>
<tr><td>STATI-32</td><td>FFFFh:8000h</td><td>…</td><td>…</td><td>Vendor Unique</td><td>Statistic Identifier’s between 8000h and FFFFh are vendor unique.</td></tr>
</table>

4.9.12 Statistic Descriptor Location

The Statistic Area location in Telemetry Data Area 1 can be found by using the Data Area 1 Statistic Area 1 Start and Data Area 1 Statistic Area 1 Size. This can be seen in the Data Area 1 Statistics Figure below.

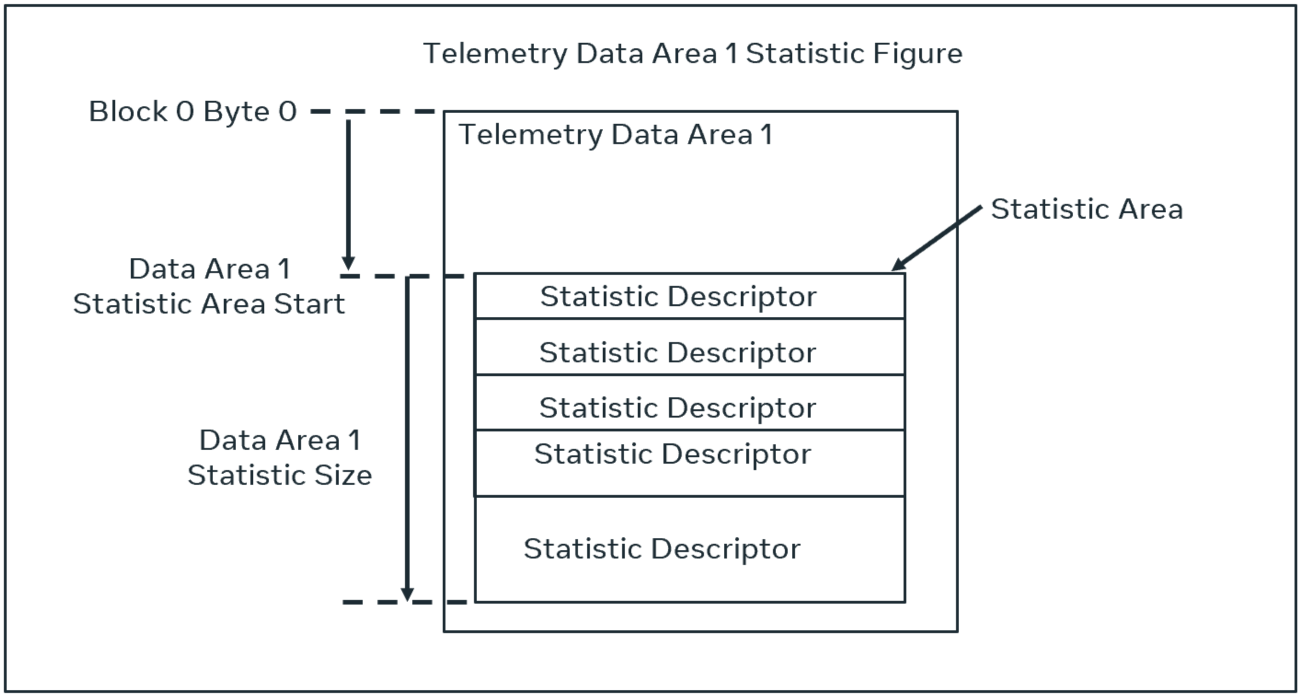

The Statistic Area location in Data Area 2 can be found by using the Data Area 2 Statistic Area Start and Data Area 2 Statistic Size. This can be seen in the Data Area 2 Statistic Figure below.

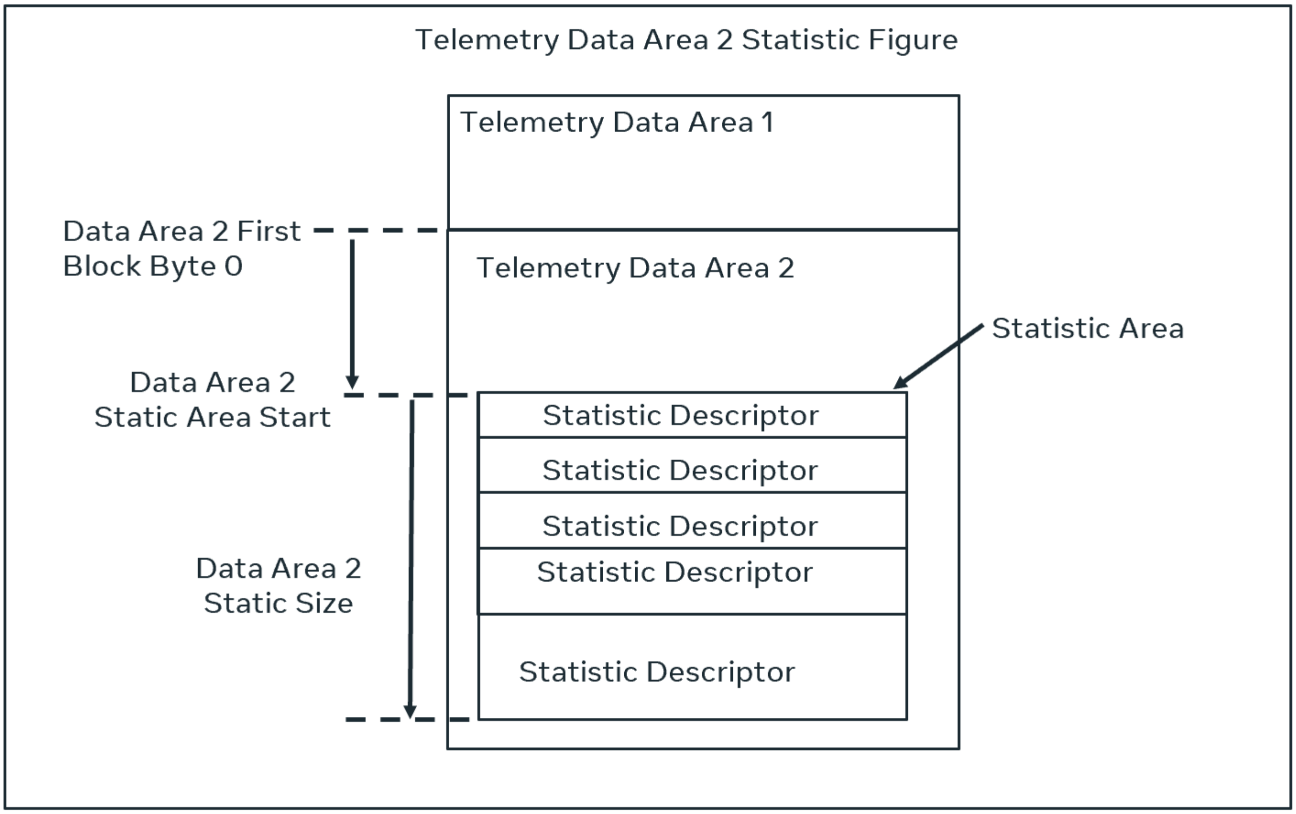

4.9.13 Statistic Requirements

Requirement ID	Description
SDL-1	Statistic Area 1 shall not cross from Telemetry Data Area 1 to Telemetry Data Area 2.
SDL-2	Bytes in the Statistics Areas that are not contained in a Statistic Descriptor shall be zero.
SDL-3	The decode of the Statistic Identifiers shall be provided to the customer via the Telemetry Strings Log.

4.9.14 Event FIFO Requirements

4.9.14.1 Event Descriptor

The Event Descriptor is used to log debug events in the Event FIFOs.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
EVENTD-1	0	Debug Event Class Type	1	Shall contain the corresponding Event Class Type for this event.
EVENTD-2	2:1	Event ID	2	Shall contain the corresponding Event ID for this event.
EVENTD-3	3	Event Data Size	1	Shall contain the number of Dwords that is used by the Event Specific Data field.
EVENTD-4	(Event Data Size * 4) + 3:4	Event Specific Data	Event Data Size *4	Shall contain event specific extended information for this event (if any).

4.9.14.2 Event Classes

This section defines the Event Class information that can be used to debug issues with the device.

4.9.14.2.1 Debug Event Class Types

Debug Event Class Code	Description
00h	Reserved
01h	Timestamp Debug Class
02h	PCIe Debug Class
03h	NVMe Debug Class
04h	Reset Debug Class
05h	Boot Sequence Debug Class
06h	Firmware Assert Debug Class
07h	Temperature Debug Class
08h	Media Debug Class
09h	Media Wear Class
0Ah	Statistic Snapshot Class

Debug Event Class Code	Description
7Fh-0Bh	Reserved
FFh-80h	Vendor Unique Class

4.9.14.2.2 Timestamp Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-TS-1</td><td>0</td><td>Timestamp Debug Class</td><td>1</td><td>Shall be set to 01h.</td></tr>
<tr><td>EVC-TS-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>Timestamp Host Command Issued</td></tr>
<tr><td>0001h</td><td>Timestamp Snapshot</td></tr>
<tr><td>0002h</td><td>Timestamp is Power on Hours</td></tr>
<tr><td>7FFFh-0003h</td><td>Reserved</td></tr>
<tr><td>FFFFh-8000h</td><td>Vendor Unique</td></tr>
</table>
</td></tr>
<tr><td>EVC-TS-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 2h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 2h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-TS-4</td><td>11:4</td><td>Timestamp</td><td>8</td><td>Time stamp value based on Event ID, in the format defined in NVM Express</td></tr>
<tr><td>EVC-TS-5</td><td>13:12</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-TS-6</td><td>(Event Data Size * 4) + 11:14</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.3 PCIe Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-PCI-1</td><td>0</td><td>PCIe Debug Class</td><td>1</td><td>Shall be set to 02h.</td></tr>
<tr><td>EVC-PCI-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log PCIe Events as follows:
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>Link Up.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>0001h</td><td>Link Down.</td></tr>
<tr><td>0002h</td><td>PCIe Error Detected.</td></tr>
<tr><td>0003h</td><td>PERST# Asserted.</td></tr>
<tr><td>0004h</td><td>PERST# De-asserted.</td></tr>
<tr><td>0005h</td><td>Refclk Stable.</td></tr>
<tr><td>0006h</td><td>Vmain Stable.</td></tr>
<tr><td>0007h</td><td>Link Speed and Width Negotiated.</td></tr>
<tr><td>7FFFh-0008h</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table></td></tr>
<tr><td>EVC-PCI-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 1h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 1h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-PCI-4</td><td>7:4</td><td>PCIe Debug Event Data</td><td>4</td><td>For PCIe Debug Event ID 7, the device shall log the negotiated Link Speed and Width as follows:<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>4</td><td>State changed flags:
● 00h = Unchanged
● 01h = Link Speed Changed
● 02h = Link Width Changed
● 03h – FFh = Reserved</td></tr>
<tr><td>5</td><td>Link Speed:
● 00h = Reserved
● 01h = PCIe Gen1
● 02h = PCIe Gen2
● 03h = PCIe Gen3
● 04h = PCIe Gen4
● 05h = PCIe Gen5
● 06h = PCIe Gen6
● 07h = PCIe Gen7
● 08h – FFh = Reserved</td></tr>
<tr><td>6</td><td>Link Width:
● 00h = Reserved
● 01h = x1</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td></td><td>● 02h = x2
● 03h = x4
● 04h = x8
● 05h = x16
● 06h – FFh = Reserved</td></tr><tr><td>7</td><td>Reserved. Shall be cleared to zero.</td></tr></table>For all other PCIe Debug Event IDs, this field shall be cleared to zero.</td></tr>
<tr><td>EVC-PCI-5</td><td>9:8</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-PCI-6</td><td>(Event Data Size * 4) + 7:10</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.4 NVMe Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-NVME-1</td><td>0</td><td>NVMe Debug Class</td><td>1</td><td>Shall be set to 03h.</td></tr>
<tr><td>EVC-NVME-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log NVMe Events as follows:<table><tr><th>Event ID</th><th>Description</th></tr><tr><td>0000h</td><td>CC.EN transitions from 0b to 1b.</td></tr><tr><td>0001h</td><td>CC.EN transitions from 1b to 0b.</td></tr><tr><td>0002h</td><td>CSTS.RDY transitions from 0b to 1b.</td></tr><tr><td>0003h</td><td>CSTS.RDY transitions from 1b to 0b.</td></tr><tr><td>0004h</td><td>Reserved</td></tr><tr><td>0005h</td><td>Create I/O Submission Queue Command or Create I/O Completion Queue Command Processed.</td></tr><tr><td>0006h</td><td>Other Admin Queue Command Processed.</td></tr><tr><td>0007h</td><td>An Admin Command Returned a Non-zero Status Code.</td></tr><tr><td>0008h</td><td>An I/O Command Returned a Non-zero Status Code.</td></tr><tr><td>0009h</td><td>CSTS.CFS Transitioned from 0b to 1b.</td></tr></table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><td>000Ah</td><td>Admin Submission Queue Base Address Written (AQA) or Admi Completion Queue Based Address (ACQ) written</td></tr>
<tr><td>000B</td><td>Controller Configuration Register (CC) Changed except for the cases that are covered in 0000h and 0001h.</td></tr>
<tr><td>000C</td><td>Controller Status Register (CSTS) Changed except for the cases that are covered in 0002h and 0003h.</td></tr>
<tr><td>000D</td><td>Delete I/O Completion Queue Command or Delete I/O Submission Queue Command Processed.</td></tr>
<tr><td>7FFFh-000Eh</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table>
</td></tr>
<tr><td>EVC-NVME-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 2h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 2h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-NVME-4</td><td>11:4</td><td>NVMe Debug Event Data</td><td>8</td><td>For NVMe Debug Event IDs 07h and 08h, the device shall log the NVMe Debug Event Data as follows:
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>4</td><td>Command Opcode</td></tr>
<tr><td>6:5</td><td>Status Code in bits 14:0 and bit 15 shall be cleared to 0b.</td></tr>
<tr><td>8:7</td><td>Command Identifier (CID)</td></tr>
<tr><td>10:9</td><td>Submission Queue Identifier (SQID)</td></tr>
<tr><td>11</td><td>Reserved. Shall be cleared to 0h.</td></tr>
</table>
For NVMe Debug Event ID 0Bh the device shall log the NVMe Debug Event Data as follows:</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>7:4</td><td>Controller Configuration Register</td></tr>
<tr><td>11:8</td><td>Reserved. Shall be cleared to 0h.</td></tr>
</table>
For NVMe Debug Event ID 0Ch the device shall log the NVMe Debug Event Data as follows:
<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>7:4</td><td>Controller Status Register</td></tr>
<tr><td>11:8</td><td>Reserved. Shall be cleared to 0h.</td></tr>
</table>
For all other NVMe Debug Event IDs, this field shall be cleared to zero.</td></tr>
<tr><td>EVC-NVME-5</td><td>13:12</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-NVME-6</td><td>(Event Data Size * 4) + 11:14</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.5 Reset Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-RST-1</td><td>0</td><td>Reset Debug Class</td><td>1</td><td>Shall be set to 04h.</td></tr>
<tr><td>EVC-RST-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log Reset Events as follows:
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>PCIe Conventional Hot Reset.</td></tr>
<tr><td>0001h</td><td>Main Power Cycle.</td></tr>
<tr><td>0002h</td><td>PERST#.</td></tr>
<tr><td>0003h</td><td>PCIe Function Level Reset.</td></tr>
<tr><td>0004h</td><td>NVM Subsystem Reset.</td></tr>
<tr><td>7FFFh-0005h</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table></td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
EVC-RST-3	3	Event Data Size	1	Event Data Size in Dwords shall be 0h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 0h. If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.
EVC-RST-4	5:4	VU Event Identifier	2	Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.
EVC-RST-5	(Event Data Size * 4) + 3:6	VU Data	(VU Data Size * 4) – 2	Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.

4.9.14.2.6 Boot Sequence Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-BS-1</td><td>0</td><td>Boot Sequence Debug Class</td><td>1</td><td>Shall be set to 05h.</td></tr>
<tr><td>EVC-BS-2</td><td>2:1</td><td>Event ID</td><td>1</td><td>The device shall log Boot Sequence Events as follows:
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>Main Firmware Boot Complete. For a single controller subsystem, this is when all boot activities required to allow I/O commands to be processed are completed excluding the enablement of the controller. This is always after the FTL load from NVM Complete, FTL Rebuild Started and FTL Rebuild Completes events.</td></tr>
<tr><td>0001h</td><td>FTL Load from NVM Complete. This means the SSD downloaded the entire FTL table from NVM to internal memory (i.e., no journaling data is required to be loaded).</td></tr>
<tr><td>0002h</td><td>FTL Rebuild Started. This is after FTL Load from NVM has completed, if the FTL table is required to be rebuilt from journaling data, then mark the start.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>0003h</td><td>FTL Rebuild Complete. This means the FTL table is ready to fully support I/O commands.</td></tr>
<tr><td>7FFFh-0004h</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table></td></tr>
<tr><td>EVC-BS-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 0h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 0h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-BS-4</td><td>5:4</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-BS-5</td><td>(Event Data Size * 4) + 3:6</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.7 Firmware Assert Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-ASRT-1</td><td>0</td><td>Firmware Assert Debug Class</td><td>1</td><td>Shall be set to 06h.</td></tr>
<tr><td>EVC-ASRT-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log Firmware Assert Events as follows:<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>Assert in NVMe Processing Code.</td></tr>
<tr><td>0001h</td><td>Assert in Media Code.</td></tr>
<tr><td>0002h</td><td>Assert in Security Code.</td></tr>
<tr><td>0003h</td><td>Assert in Background Services Code.</td></tr>
<tr><td>0004h</td><td>FTL Rebuild Failed.</td></tr>
<tr><td>0005h</td><td>FTL Data Mismatch.</td></tr>
<tr><td>0006h</td><td>Assert in Other Code.</td></tr>
<tr><td>7FFFh-0007h</td><td>Reserved</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr></table></td></tr>
<tr><td>EVC-ASRT-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 0h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 0h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-ASRT-4</td><td>5:4</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-ASRT-5</td><td>(Event Data Size * 4) + 3:6</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.8 Temperature Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-TEMP-1</td><td>0</td><td>Temperature Debug Class</td><td>1</td><td>Shall be set to 07h.</td></tr>
<tr><td>EVC-TEMP-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log Temperature Events as follows:<table><tr><th>Event ID</th><th>Description</th></tr><tr><td>0000h</td><td>Composite Temperature decreases to (WCTEMP – 2)</td></tr><tr><td>0001h</td><td>Composite Temperature increases to WCTEMP</td></tr><tr><td>0002h</td><td>Composite Temperature increases to reach CCTEMP.</td></tr><tr><td>7FFFh-0003h</td><td>Reserved</td></tr><tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr></table>Once the event is logged 2 degrees of hysteresis should be applied to avoid additional entries being generated.</td></tr>
<tr><td>EVC-TEMP-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 0h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 0h.</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
				If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.
EVC-TEMP-4	5:4	VU Event Identifier	2	Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.
EVC-TEMP-5	(Event Data Size * 4) + 3:6	VU Data	(VU Data Size * 4) – 2	Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.

4.9.14.2.9 Media Debug Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-MEDIA-1</td><td>0</td><td>Media Debug Class</td><td>1</td><td>Shall be set to 08h.</td></tr>
<tr><td>EVC-MEDIA-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log Media Events as follows:
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>XOR (or equivalent) Recovery Invoked.</td></tr>
<tr><td>0001h</td><td>Uncorrectable Media Error</td></tr>
<tr><td>0002h</td><td>Block Marked Bad Due to Program Error.</td></tr>
<tr><td>0003h</td><td>Block Marked Bad Due to Erase Error.</td></tr>
<tr><td>0004h</td><td>Block Marked Bad Due to Read Error.</td></tr>
<tr><td>0005h</td><td>Plane Failure Event.</td></tr>
<tr><td>7FFFh-0006h</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table>
</td></tr>
<tr><td>EVC-MEDIA-3</td><td>3</td><td>Event Data Size</td><td>0</td><td>Event Data Size in Dwords shall be 0h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 0h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-MEDIA-4</td><td>5:4</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
EVC-MEDIA-5	(Event Data Size * 4) + 3:6	VU Data	(VU Data Size * 4) – 2	Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.

4.9.14.2.10 Media Wear Event Class Format

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-WEAR-1</td><td>0</td><td>Media Wear Class</td><td>1</td><td>Shall be set to 09h. This event shall be logged once every 24 hours.</td></tr>
<tr><td>EVC-WEAR-2</td><td>2:1</td><td>Event Identifier</td><td>2</td><td>The device shall log Wear Events as follows:
<table>
<tr><th>Event ID</th><th>Description</th></tr>
<tr><td>0000h</td><td>Media Wear</td></tr>
<tr><td>7FFFh-0001h</td><td>Reserved</td></tr>
<tr><td>8000h-FFFFh</td><td>Vendor Unique</td></tr>
</table></td></tr>
<tr><td>EVC-WEAR-3</td><td>3</td><td>Event Data Size</td><td>1</td><td>Event Data Size in Dwords shall be 3h + VU Event Identifier + VU Data Size in Dwords. If VU Event Identifier and VU data do not exist, then this is 3h.
If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.</td></tr>
<tr><td>EVC-WEAR-4</td><td>15:4</td><td>Current Media Wear</td><td>12</td><td>For Wear Debug Event ID 00h, the device shall log the current wear as follows:
<table>
<tr><th>Byte Address</th><th>Description</th></tr>
<tr><td>7:4</td><td>Host Terabytes Written.</td></tr>
<tr><td>11:8</td><td>Media Terabytes Written.</td></tr>
<tr><td>15:12</td><td>Media Terabytes Erased.</td></tr>
</table>
For all other Wear Event IDs, this field shall be cleared to zero.</td></tr>
<tr><td>EVC-WEAR-4</td><td>17:16</td><td>VU Event Identifier</td><td>2</td><td>Vendor Unique Event Identifier. The VU Event Identifier that enables the VU data to be decoded.</td></tr>
<tr><td>EVC-WEAR-5</td><td>(Event Data Size * 4) + 14:18</td><td>VU Data</td><td>(VU Data Size * 4) – 2</td><td>Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.</td></tr>
</table>

4.9.14.2.11 Statistic Snapshot Class

For the Statistic Snapshot Class. This enables the device to take a snapshot of a statistic descriptor and store this value in the Event FIFO.

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>EVC-SNP-1</td><td>0</td><td>Statistic Snapshot Class</td><td>1</td><td>Shall be set to 0Ah. This is a statistic snapshot of a statistics descriptor.</td></tr>
<tr><td>ENC-SNP-2</td><td>3:1</td><td>Reserved</td><td>3</td><td>Reserved. Cleared to 0h.</td></tr>
<tr><td>ENC-SNP-3</td><td>5:4</td><td>Statistic Identifier</td><td>2</td><td>Shall contain the identifier that defines the data format of this entry.</td></tr>
<tr><td>ENC-SNP-4</td><td>6</td><td>Statistic Info</td><td>1</td><td><table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:4</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>3:0</td><td>Behavior Type</td></tr>
</table></td></tr>
<tr><td>ENC-SNP-5</td><td>7</td><td>NS INFO</td><td>1</td><td><table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7</td><td>Namespace Information Valid. When set to 1b the Namespace Identifier is valid. When cleared to 0b the Namespace Identifier is not valid</td></tr>
<tr><td>6:0</td><td>Namespace Identifier. When Namespace Information valid is set to 1b this is the Namespace Identifier (NS ID). When Namespace Information is cleared to 0b this field is not valid and shall be cleared to 0h.</td></tr>
</table></td></tr>
<tr><td>ENC-SNP-6</td><td>9:8</td><td>Statistic Data Size</td><td>2</td><td>Shall contain the number of Dwords that are used by the Statistic Specific Data field. Valid values for this field are 0h to FFFFh.</td></tr>
<tr><td>ENC-SNP-7</td><td>11:10</td><td>Reserved</td><td>2</td><td>Reserved. Shall be cleared to 0h.</td></tr>
<tr><td>ENC-SNP-8</td><td>(Statistic Data Size * 4) + 11:12</td><td>Statistic Specific Data</td><td>Statistic Data Size * 4</td><td>Shall contain data specific extended information for this Statistic Data Entry type if this exists.</td></tr>
</table>

4.9.14.2.12 Vendor Unique Event Class Format

For the Vendor Unique debug class.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
EVC-VU-1	0	Debug VU Class	1	Shall be set to 80-FFh.
EVC-VU-2	2:1	VU Event Identifier	2	Vendor Unique Event Identifier. Allowable Vendor Unique event identifiers under the debug VU Class is 0h-FFFFh.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
EVC-VU-3	3	VU Event Data Size	1	VU Event Data Size in Dwords shall be the VU Data Size in Dwords. If there is no VU Data, then the VU Event Identifier field and VU Data field do not exist (i.e., the VU Data Size is 0h) and the Event Data Size shall be set to 1h.
EVC-VU-4	(Event Data Size * 4) + 3:4	VU Data	VU Data Size * 4	Vendor Unique Data. The size of this data is determined based on the Event Data Size if this field exists.

4.9.15 Event FIFO Location

The Event FIFO X location in Telemetry Data Area 1 can be found by using the Event FIFO X Start, Event FIFO X Size. This can be seen in the Telemetry Data Area 1 Event FIFO Figure below.

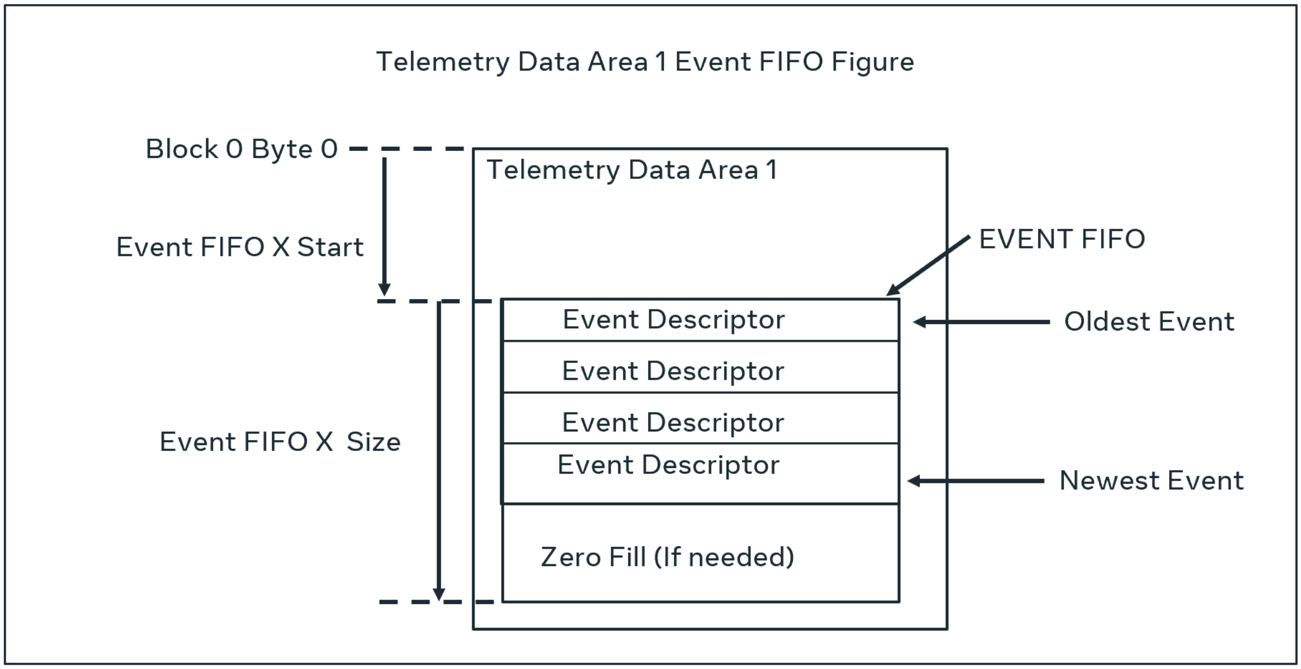

The Event FIFO location in Telemetry Data Area 2 can be found by using the Event FIFO X Start, Event FIFO X Data End. This can be seen in the Telemetry Data Area 2 Event FIFO Figure below.

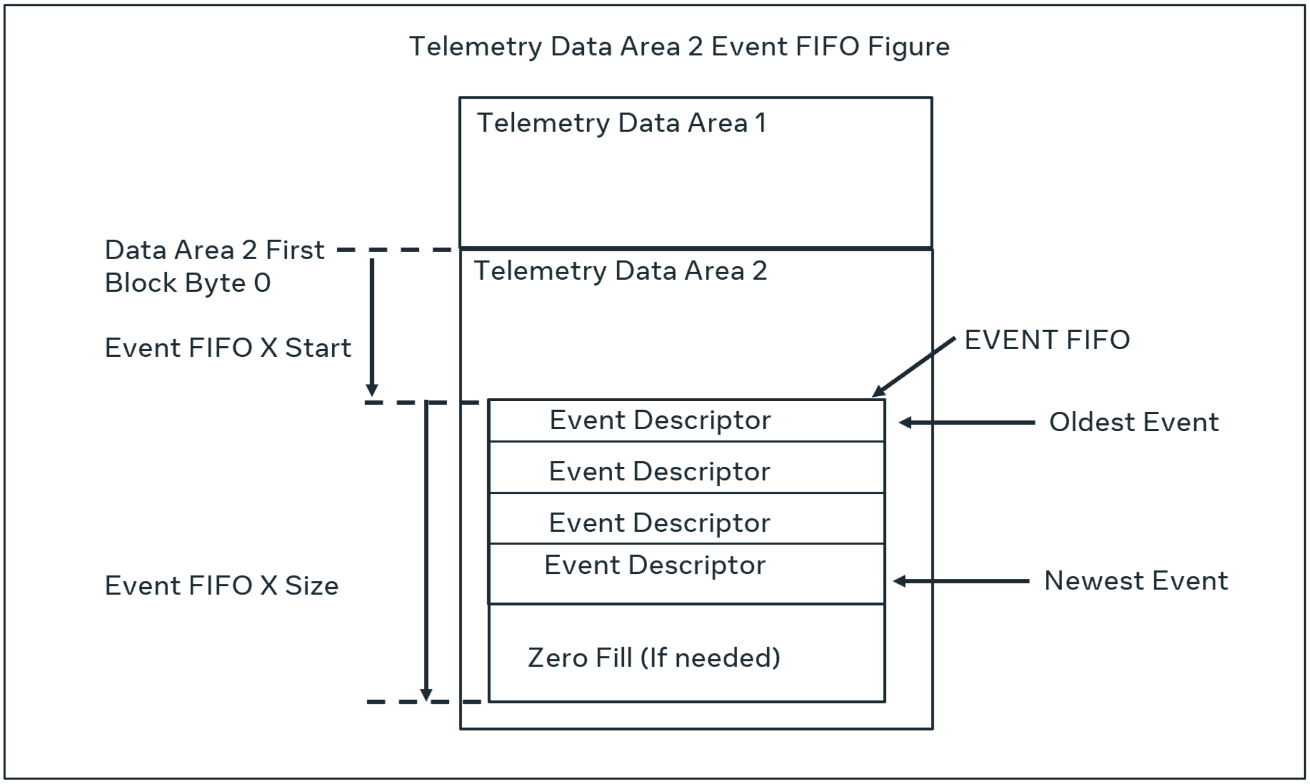

4.9.15.1 Event FIFO Requirements

Requirement ID	Description
EVF-1	The oldest event descriptor shall be at the Event FIFO Start Address.
EVF-2	There shall be no gaps between the data in one descriptor and the next descriptor.
EVF-3	The event FIFO shall only have complete descriptors (i.e., no partial descriptors).
EVF-4	All data contained in the event FIFO that is not contained in an event descriptor shall be cleared to zero.
EVF-5	For a single FIFO, the Event FIFO Start and Event FIFO End shall be in the same Telemetry Data Area.
EVF-6	Event FIFO Start and End locations shall not overlap with the Statistic Descriptor locations.
EVF-7	The decode of the Event Class and Vendor Unique Event Class shall be provided to the customer via the Telemetry String Log page.

4.9.16 Telemetry Data Area

4.9.16.1 Telemetry Data Format

Telemetry Data Area 1 has the following format. It should be noted that this is the format of Telemetry Data Area 1 in the Telemetry log, which does not start from byte 0/Dword 0.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELDF-1	1:0	Major Version	2	This field shall be set to 0003h. A change of the Major Version number indicates that the changes to the Telemetry Data Area format decode are not host software backward compatible.
TELDF-2	3:2	Minor Version	2	This field shall be set to 0001h. A change of the Minor Version number indicates that the changes to the Telemetry Data Area format decode are host software backward compatible.
TELDF-3	7:4	Reserved	4	Shall be cleared to zero.
TELDF-4	15:8	Timestamp	8	Shall be set to the Timestamp when the Telemetry Log (Host Initiated or Controller Initiated) Get Log command was received by the device. The format of this field shall be as defined in Timestamp (Feature Identifier 0Eh) of the NVM Express Specification.
TELDF-5	31:16	Log Page GUID	16	This field shall be set to BA560A9C3043424CBC73719D87E64EFAh.
TELDF-6	32	Number Telemetry Profiles Supported	1	Shall be set to the number of Telemetry Profiles supported by the device. This field is zero based.
TELDF-7	33	Telemetry Profile Selected (TPS)	1	Shall be set to the currently selected Telemetry Profile. This field is zero based.
TELDF-8	39:34	Reserved	6	Shall be cleared to zero.
TELDF-9	47:40	Telemetry String Log Size (SLS)	8	Shall be set to the number of Dwords in the Telemetry String Log.
TELDF-10	55:48	Reserved	8	Shall be cleared to zero.
TELDF-11	63:56	Firmware Revision	8	Shall be set to the firmware revision reported in the Identify Controller Data Structure.
TELDF-12	95:64	Reserved	32	Shall be cleared to zero.
TELDF-13	103:96	Data Area 1 Statistic Start	8	This field shall contain the Dword offset of the start of the Statistic Buffer in Telemetry Data Area 1. This offset starts from Dword 0 / byte 0 in Telemetry Data Area 1. This is a zero based value.
TELDF-14	111:104	Data Area 1 Statistic Size	8	This field shall contain the size of the Statistic Buffer in Data Area 1 in Dwords. If there are no Statistic Descriptors a value of 0h shall be reported.
TELDF-15	119:112	Data Area 2 Statistic Start	8	This field shall contain the Dword offset of the start of the Statistic Buffer in Telemetry Data Area 2. This

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td>offset starts from Dword 0 / byte 0 in Telemetry Data Area 2. This is a zero based value.</td></tr>
<tr><td>TELDF-16</td><td>127:120</td><td>Data Area 2 Statistic Size</td><td>8</td><td>This field shall contain the size of the Statistic Buffer in Data Area 2 in Dwords. If there are no Statistic Descriptors a value of 0h shall be reported.</td></tr>
<tr><td>TELDF-17</td><td>159:128</td><td>Reserved</td><td>32</td><td>Shall be cleared to zero.</td></tr>
<tr><td>TELDF-18</td><td>160</td><td>Event FIFO 1 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 1 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 1 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 1 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-19</td><td>161</td><td>Event FIFO 2 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 2 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 2 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 2 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-20</td><td>162</td><td>Event FIFO 3 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 3 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 3 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 3 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-21</td><td>163</td><td>Event FIFO 4 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 4 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 4 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 4 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>TELDF-22</td><td>164</td><td>Event FIFO 5 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 5 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 5 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 5 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-23</td><td>165</td><td>Event FIFO 6 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 6 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 6 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 6 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-24</td><td>166</td><td>Event FIFO 7 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 7 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 7 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 7 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-25</td><td>167</td><td>Event FIFO 8 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 8 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 8 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 8 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-26</td><td>168</td><td>Event FIFO 9 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 9 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 9 exists in Telemetry Data Area 1</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table><tr><td>02h</td><td>Event FIFO 9 exists in Telemetry Data Area 2</td></tr><tr><td>0F-03h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>TELDF-27</td><td>169</td><td>Event FIFO 10 Data Area</td><td>1</td><td><table><tr><th>Event FIFO Data Area</th><th>Description</th></tr><tr><td>00h</td><td>Event FIFO 10 does not exist</td></tr><tr><td>01h</td><td>Event FIFO 10 exists in Telemetry Data Area 1</td></tr><tr><td>02h</td><td>Event FIFO 10 exists in Telemetry Data Area 2</td></tr><tr><td>0F-03h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>TELDF-28</td><td>170</td><td>Event FIFO 11 Data Area</td><td>1</td><td><table><tr><th>Event FIFO Data Area</th><th>Description</th></tr><tr><td>00h</td><td>Event FIFO 11 does not exist</td></tr><tr><td>01h</td><td>Event FIFO 11 exists in Telemetry Data Area 1</td></tr><tr><td>02h</td><td>Event FIFO 11 exists in Telemetry Data Area 2</td></tr><tr><td>0F-03h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>TELDF-29</td><td>171</td><td>Event FIFO 12 Data Area</td><td>1</td><td><table><tr><th>Event FIFO Data Area</th><th>Description</th></tr><tr><td>00h</td><td>Event FIFO 12 does not exist</td></tr><tr><td>01h</td><td>Event FIFO 12 exists in Telemetry Data Area 1</td></tr><tr><td>02h</td><td>Event FIFO 12 exists in Telemetry Data Area 2</td></tr><tr><td>0F-03h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>TELDF-30</td><td>172</td><td>Event FIFO 13 Data Area</td><td>1</td><td><table><tr><th>Event FIFO Data Area</th><th>Description</th></tr><tr><td>00h</td><td>Event FIFO 13 does not exist</td></tr><tr><td>01h</td><td>Event FIFO 13 exists in Telemetry Data Area 1</td></tr><tr><td>02h</td><td>Event FIFO 13 exists in Telemetry Data Area 2</td></tr><tr><td>0F-03h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>TELDF-31</td><td>173</td><td>Event FIFO 14 Data Area</td><td>1</td><td><table><tr><th>Event FIFO Data Area</th><th>Description</th></tr><tr><td>00h</td><td>Event FIFO 14 does not exist</td></tr></table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>01h</td><td>Event FIFO 14 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 14 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-32</td><td>174</td><td>Event FIFO 15 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 15 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 15 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 15 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-33</td><td>175</td><td>Event FIFO 16 Data Area</td><td>1</td><td><table>
<tr><th>Event FIFO Data Area</th><th>Description</th></tr>
<tr><td>00h</td><td>Event FIFO 16 does not exist</td></tr>
<tr><td>01h</td><td>Event FIFO 16 exists in Telemetry Data Area 1</td></tr>
<tr><td>02h</td><td>Event FIFO 16 exists in Telemetry Data Area 2</td></tr>
<tr><td>0F-03h</td><td>Reserved</td></tr>
</table></td></tr>
<tr><td>TELDF-34</td><td>183:176</td><td>Event FIFO 1 Start</td><td>8</td><td>Shall be set to the Dword offset of Event FIFO 1 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.</td></tr>
<tr><td>TELDF-35</td><td>191:184</td><td>Event FIFO 1 Size</td><td>8</td><td>This field shall contain the size of Event FIFO 1 in Dwords.</td></tr>
<tr><td>TELDF-36</td><td>199:192</td><td>Event FIFO 2 Start</td><td>8</td><td>Shall be set to the Dword offset of Event FIFO 2 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.</td></tr>
<tr><td>TELDF-37</td><td>207:200</td><td>Event FIFO 2 Size</td><td>8</td><td>This field shall contain the size of Event FIFO 2 in Dwords.</td></tr>
<tr><td>TELDF-38</td><td>215:208</td><td>Event FIFO 3 Start</td><td>8</td><td>Shall be set to the Dword offset of Event FIFO 3 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.</td></tr>
<tr><td>TELDF-39</td><td>223:216</td><td>Event FIFO 3 Size</td><td>8</td><td>This field shall contain the size of Event FIFO 3 in Dwords.</td></tr>
<tr><td>TELDF-40</td><td>231:224</td><td>Event FIFO 4 Start</td><td>8</td><td>Shall be set to the Dword offset of Event FIFO 4 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.</td></tr>
</table>

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELDF-41	239:232	Event FIFO 4 Size	8	This field shall contain the size of Event FIFO 4 in Dwords.
TELDF-42	247:240	Event FIFO 5 Start	8	Shall be set to the Dword offset of Event FIFO 5 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-43	255:248	Event FIFO 5 Size	8	This field shall contain the size of Event FIFO 5 in Dwords
TELDF-44	263:256	Event FIFO 6 Start	8	Shall be set to the Dword offset of Event FIFO 6 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-45	271:264	Event FIFO 6 Size	8	This field shall contain the size of Event FIFO 6 in Dwords.
TELDF-46	279:272	Event FIFO 7 Start	8	Shall be set to the Dword offset of Event FIFO 7 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-47	287:280	Event FIFO 7 Size	8	This field shall contain the size of Event FIFO 7 in Dwords.
TELDF-48	295:288	Event FIFO 8 Start	8	Shall be set to the Dword offset of Event FIFO 8 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-49	303:296	Event FIFO 8 Size	8	This field shall contain the size of Event FIFO 8 in Dwords.
TELDF-50	311:304	Event FIFO 9 Start	8	Shall be set to the Dword offset of Event FIFO 9 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-51	319:312	Event FIFO 9 Size	8	This field shall contain the size of Event FIFO 9 in Dwords.
TELDF-52	327:320	Event FIFO 10 Start	8	Shall be set to the Dword offset of Event FIFO 10 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-53	335:328	Event FIFO 10 Size	8	This field shall contain the size of Event FIFO 10 in Dwords.
TELDF-54	343:336	Event FIFO 11 Start	8	Shall be set to the Dword offset of Event FIFO 11 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-55	351:344	Event FIFO 11 Size	8	This field shall contain the size of Event FIFO 11 in Dwords.
TELDF-56	359:352	Event FIFO 12 Start	8	Shall be set to the Dword offset of Event FIFO 12 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.

Requirement ID	Byte Address	Field	# of Bytes	Field Description
TELDF-57	367:360	Event FIFO 12 Size	8	This field shall contain the size of Event FIFO 12 in Dwords.
TELDF-58	375:368	Event FIFO 13 Start	8	Shall be set to the Dword offset of Event FIFO 13 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-59	383:376	Event FIFO 13 Size	8	This field shall contain the size of Event FIFO 13 in Dwords.
TELDF-60	391:384	Event FIFO 14 Start	8	Shall be set to the Dword offset of Event FIFO 14 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-61	399:392	Event FIFO 14 Size	8	This field shall contain the size of Event FIFO 14 in Dwords.
TELDF-62	407:400	Event FIFO 15 Start	8	Shall be set to the Dword offset of Event FIFO 15 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-63	415:408	Event FIFO 15 Size	8	This field shall contain the size of Event FIFO 15 in Dwords.
TELDF-64	423:416	Event FIFO 16 Start	8	Shall be set to the Dword offset of Event FIFO 16 starting from Dword 0 / byte 0 of the selected Telemetry Data Area.
TELDF-65	431:424	Event FIFO 16 Size	8	This field shall contain the size of Event FIFO 16 in Dwords.
TELDF-66	511:432	Reserved	80	Shall be cleared to zero.
TELDF-67	1023:512	SMART / Health Information	512	This shall be a copy of the SMART / Health Information (Log Identifier 02h) data of the controller that captured the Telemetry data.
TELDF-68	1535:1024	SMART / Health Information Extended	512	This shall be a copy of the SMART / Health Information Extended (Log Identifier C0h) data of the controller that captured the Telemetry data.

4.10 Device Self-test Requirements

To normalize the reporting of Device Self-test data, the following Segment Numbers are defined.

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>SELFTST-1</td><td>The device shall use the following segment numbers:
<table>
<tr><th>Segment</th><th>Test Performed</th><th>Failure Criteria</th></tr>
<tr><td>0</td><td>Reserved.</td><td>N/A</td></tr>
</table>
</td></tr>
</table>

Requirement ID	Description		
	1 – SRAM Check	Write a test pattern to SRAM followed by a read and compare of original data. Shall test all internal and external SRAM devices.	Any uncorrectable error or data miscompare.
	2 – SMART Check	Check SMART or health status for Critical Warning bits set to 1b in SMART / Health Information Log.	Any Critical Warning bit set to 1b fails this segment, with the exception of bit 0 (spare capacity below threshold) and bit 1 (thermal excursions).
	3 – Volatile Memory Backup	Validate volatile memory backup solution health. Measure backup power source charge and/or discharge time. Shall perform backup test and not rely on previous results. During the test, the device shall not lose power loss protection.	Significant degradation in backup capability.
	4 – Non-User Data Validation	Confirm/validate all copies of non-user data (e.g., FTL, log pages, etc.).	Non-user data is corrupt and is not recoverable.
	5 – NVM Integrity	Write/read/compare to reserved areas of each NVM, if present. Ensure that every read/write channel of the controller is exercised and each NVM is accessed.	Data miscompare or hardware failure.
	6 – Data Integrity	Perform background housekeeping tasks, prioritizing actions that enhance the integrity of stored data. Test CRC hardware, error correction hardware, including XOR, encryption engine, etc.	Any uncorrectable data error or hardware failure.
	7 – Media Check Extended	Perform reads from every mapped LBA. Shall include LBA 0 and MAXLBA, if mapped. Shall terminate segment when all mapped LBAs are read or the Device Self-test completion time limit expires.	Any uncorrectable error.
	8 – DRAM check	Write a test pattern to DRAM followed by a read and compare of original data. Shall test all internal and external SRAM devices.	Any uncorrectable error or data miscompare.

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td><table>
<tr><td>9 – Temperature Sensor</td><td>Verify the temperature sensors are working correctly.</td><td>Temperature sensors are not reporting correctly.</td></tr>
<tr><td>10 – SMART Check</td><td>Same as 2-SMART Check. This check shall be performed after all other test segments in order to detect critical warnings caused by prior test segments.</td><td></td></tr>
<tr><td>200:11</td><td>Reserved.</td><td>N/A</td></tr>
<tr><td>255:201</td><td>Vendor Unique. Any Vendor Unique Segment Numbers shall be clearly documented and disclosed to the customer.</td><td></td></tr>
</table></td></tr>
<tr><td>SELFTST-2</td><td>A short device self-test operation (i.e., the SCT field is set to 1h) shall complete in less than or equal to 2 minutes.</td></tr>
<tr><td>SELFTST-3</td><td>An extended device self-test operation (i.e., the SCT field is set to 2h) shall complete in less than or equal to 30 minutes (i.e., the EDSTT field shall indicate a value of less than or equal to 30 minutes).</td></tr>
<tr><td>SELFTST-4</td><td>A vendor-specific self-test operation (i.e., the SCT field is set to Eh) shall complete the test specified by SELFTST-7 in less than or equal to 2 seconds.</td></tr>
<tr><td>SELFTST-5</td><td>A device self-test operation shall not fail if it is unable to complete within the specified completion time (e.g., due to concurrent host I/O activity). The device shall complete as many device self-test segments as possible within the allotted completion time.</td></tr>
<tr><td>SELFTST-6</td><td>For a Self-test Code (STC) of 2h, the device self-test operation shall terminate upon reading all allocated LBAs and completing all non-media related tests.</td></tr>
<tr><td>SELFTST-7</td><td>For a Device Self-test Code (STC) of 1h or Eh, the device self-test operation shall include at a minimum, segments 1 to 3 and 8, and as much as possible of segments 4 to 7, within the completion time specified for that specific operation.</td></tr>
<tr><td>SELFTST-8</td><td>For the SMART Check (segments 2 and 10), exceeding a temperature threshold shall not cause the device self-test operation to fail since that is generally caused by a fan/system failure. The SMART / Health Information log page Critical Warning bit 1 shall be ignored for purposes of the device self-test operation.</td></tr>
<tr><td>SELFTST-9</td><td>The device vendor shall provide the minimum number of LBAs required to be written by the host to enable the extended device self-test operation to read from all physical non-volatile memory elements (e.g., NAND die).</td></tr>
<tr><td>SELFTST-10</td><td>If the Device Self-test command is reissued after a previous Device Self-test command and operation completed successfully, then the Media Check Extended (segment 7) shall test different LBAs, wrapping around if necessary. For example, if the device contains 10000h allocated LBAs, but only 6000h LBAs are able to be scanned within the specified Device Self-test completion time, then the first run of the Device Self-test would scan LBAs 0h to 5FFFh, the next run would scan LBAs 6000h to BFFFh and the next would scan LBAs beginning at C000h wrapping to 0h and then to 1FFFh, and so forth.</td></tr>
</table>

Requirement ID	Description
SELFTST-11	The device self-test operation shall not be destructive to any user data (note that reading user data once during the device self-test operation is not considered destructive, because even though it technically causes read disturb, the effect is minimal).

4.11 Firmware Update Requirements

This defines the requirements for firmware update in the device.

Requirement ID	Description
FWUP-1	A firmware activation shall be recorded on the device.
FWUP-2	Other than media end-of-life, devices shall not have any restrictions on the number of firmware downloads supported.
FWUP-3	The Firmware Commit command with the following Commit Action (CA) codes shall be supported: • 000b – Download only. • 001b – Download and activate upon reset. • 010b – Activate upon reset. • 011b – Activate immediately without reset.
FWUP-4	Firmware Image Download command shall be supported and shall not cause any performance impact.
FWUP-5	Obsolete.
FWUP-6	The device shall support a minimum of 2 read/write slots for firmware update and may support up to 7.
FWUP-7	For firmware commit action 011b (firmware activation without reset), the device shall complete the firmware activation process and be ready to accept host I/O and NVMe Admin commands within 1 second from the receipt of the Firmware Commit command. The Maximum Time for Firmware Activation (MTFA) field shall not exceed Ah.
FWUP-8	If a Firmware Commit command is processed that specifies a Security Firmware Revision (see SMART-22) that is less than the Security Firmware of the actively running firmware revision, then the device shall return a status code of Firmware Activation Prohibited (13h). If a firmware image is committed/activated that has a Security Firmware Revision that is greater than or equal to the Security Firmware Revision of the actively running firmware revision, then the operation shall be permitted.
FWUP-9	The firmware image in each valid firmware slot shall have multiple copies of that firmware image on separate physical non-volatile memory devices (e.g., NAND die, NAND RAID stripe) for reliability.
FWUP-10	Firmware activation shall not cause or require data (e.g., user data, log page contents, etc.) to be lost or destroyed.
FWUP-11	Firmware activation without reset shall occur seamlessly on a running system: • The device shall preserve the current state of the device (e.g., Opal locking state, current value of each Feature, Timestamp, I/O Submission/Completion Queue setup,

Requirement ID	Description
	NVMe-MI state such as MCTP Transmission Unit Size, SMBus/I2C Frequency, and MCTP state such as the MCTP endpoint ID, LBA format, PCIe state, etc.). • The PCIe link state shall not be impacted.
FWUP-12	The firmware image shall contain the firmware for all the device's controllers (NVMe controllers, management controllers, etc.).
FWUP-13	The VPD/FRU Data shall be a part of the device's firmware update image and shall be field upgradable using the standard NVMe Firmware Image Download/Firmware Commit commands. VPD fields that are unique to the device (e.g., Serial Number) shall be preserved during a firmware update.
FWUP-14	The size of the firmware should be minimized in order to facilitate firmware update via low-bandwidth interfaces (e.g., NVMe-MI over SMBus/PCIe VDM). It is recommended that the firmware be 6 MB or less.
FWUP-15	Requiring a reset of any type to activate firmware under any condition shall be avoided. If a reset is unavoidable, the device vendor shall request a waiver from the end customer. In cases where waivers are granted, a Controller Level Reset is preferred over a Conventional Reset. Firmware Activation requiring an NVMe Subsystem Reset should always be avoided.
FWUP-16	All firmware updates shall be allowed to be performed by the system vendor or end user without any outside assistance from NVMe device vendor employees or anyone else and shall not require any additional vendor-specific actions or equipment (such as dongles, debug cables, etc.) unless otherwise noted in this specification. This includes both pre-production and production firmware.
FWUP-17	Firmware updates shall not require an intermediate firmware update or "bridge code" to update from one firmware version to the next subsequent version (e.g., N-1 firmware revision to bridge firmware to N firmware is prohibited). This includes both pre-production and production firmware.
FWUP-18	The device shall not be left in an unusable state due to any reset or loss of power during a firmware update. Depending on where in the device's firmware activation process the reset or loss of power occurs, the device shall either: • abort the firmware activation and continue running with the firmware revision that was running at the start of the firmware update process; or • complete the firmware activation and start running with the new firmware revision.
FWUP-19	For production-level firmware, all means of updating the firmware besides the Firmware Image Download/Firmware Commit commands (e.g., via MMIO or any vendor-proprietary mechanisms) shall be disabled.
FWUP-20	The device shall support a read-only firmware slot (slot 1). If a firmware image is successfully activated to a read-write slot that increases the Security Version Number (see SMART-22), then the device shall also update the read-only slot 1 with that firmware image.
FWUP-21	If the Firmware Slot field in the Firmware Commit command is cleared to 0h, then the device shall commit the firmware image to a firmware slot that is not currently active.
FWUP-22	If the current active firmware image is corrupt, then the device shall automatically fail over to another backup copy of that firmware image until all backup copies of that firmware image are exhausted. The device shall not automatically fail over to another revision of

Requirement ID	Description
	firmware in a different firmware slot, this overrides the NVM Express Base Specification. The vendor should follow the guidelines in NIST SP 800-193.
FWUP-23	A single firmware image shall cover all the devices in a product family. For example, devices with the same controller but different capacity points, endurance or form factor shall have one firmware image that works on all devices. A single firmware image shall cover all TCG Opal capable and non-TCG Opal capable devices in a given product family. An exception is that the firmware for a family of FIPS 140-3 compliant devices is different than the firmware for the same family of non-FIPS 140-3 compliant devices. For example, the firmware for the family of FIPS 140-3 compliant devices shall contain the FIPS security descriptor whereas the firmware for the family of non-FIPS 140-3 compliant devices shall not contain the FIPS security descriptor. Other than the FIPS security descriptor, the firmware for the family of FIPS 140-3 compliant devices should be the same as the firmware for the family of non-FIPS 140-3 compliant devices.
FWUP-24	Every firmware image shall have a firmware signature that meets the requirements of SBT-1. The device shall check the firmware signature as part of processing any Firmware Commit command that commits a firmware image to a firmware slot. If the firmware signature check fails, the device shall abort the Firmware Commit command with Invalid Firmware Image status.
FWUP-25	When activating a firmware image to the same slot as the currently running firmware (where the currently running firmware is the firmware that processed the Firmware Commit command): • the device shall keep a copy of the currently running firmware until the newly activated firmware completes initialization successfully at least once; and • if the newly activated firmware cannot complete initialization successfully, then the device shall revert to the firmware that processed the Firmware Commit command; and • once the newly activated firmware completes initialization successfully at least once, the firmware image for the firmware that processed the Firmware Commit command is permitted to be discarded.

4.12 Factory Default Requirements

This section defines additional requirements for the state of the device as shipped from the supplier factory.

Requirement ID	Description
FDEF-1	All firmware slots shall contain the most recent production firmware revision.
FDEF-2	All logical blocks shall be deallocated.
FDEF-3	The following SMART / Health Information log page fields shall be cleared to 0h: Critical Warning, Endurance Group Critical Warning Summary, Data Units Read, Data Units Written, Host Read Commands, Host Write Commands, Controller Busy Time, Power Cycles, Power On Hours, Unsafe Shutdowns, Media, and Data Integrity Errors, Percentage Used, and Number of Error Information Log Entries.

Requirement ID	Description
FDEF-4	The Available Spare field in the SMART / Health Information log page shall be set to a value of 100.
FDEF-5	The Total Number of Events field in the Persistent Event log page shall be cleared to 0h.
FDEF-6	Features that are saveable shall not have a saved value unless otherwise specified.

4.13 De-Allocation Requirements

Requirement ID	Description
TRIM-1	The device shall support the Deallocate attribute of the Dataset Management command.
TRIM-2	Once the device has deallocated a logical block a read of that logical block shall return all bytes cleared to 0h until that logical block has been modified by a command (e.g., Write, Copy, etc.) (i.e., bits 2:0 = 001b in the Deallocate Logical Block Features (DLFEAT) field in the Identify Namespace data structure).
TRIM-3	Obsolete (see TRIM-2).
TRIM-4	Device deallocated logical blocks shall provide the performance and reliability benefits of overprovisioned space.
TRIM-5	The device shall support Garbage Collection during periods of no I/O (Idle GC).
TRIM-6	Read latency shall not change more than 5% from baseline when the host is issuing De-Allocate/TRIM commands.
TRIM-7	Read latency shall not change more than 5% from baseline when the device is performing Idle GC.
TRIM-8	If any device deallocated logical blocks have been subsequently read or written to, then the contents of those logical blocks shall not change as a consequence of an unsafe power down event.
TRIM-9	The device should deallocate data to the maximum extent possible provided that it does not violate TRIM-6 or TRIM-7. Background deallocate processing is allowed provided that it does not violate TRIM-8.

4.14 Sector Size and Namespace Support

Requirement ID	Description
SECTOR-1	Obsolete. Replaced by SECTOR-4.
SECTOR-2	Obsolete. Replaced by SECTOR-4.
SECTOR-3	The device shall have one active Namespace whose size is the maximum capacity as shipped from the factory.
SECTOR-4	The device shall support 512-byte and 4096-byte logical block sizes.
SECTOR-5	The device shall comply with the latest revision of the SFF-8447 Specification for LBA Count for Disk Drives.

4.15 Set/Get Features Requirements

This section describes features defined by this specification. Unless otherwise specified, the device shall support all the features defined in this section.

4.15.1 General Feature Requirements

Requirement ID	Description
GETF-1	For any Get Feature Identifier defined in this section, Selection (SEL) values 00b to 11b in Dword 10 shall be supported.
GETF-2	If the feature requested by Set Feature is not supported, then a status error code of 02h (Invalid Field in Command) shall be returned.
GETF-3	The device shall not implement a Feature with ID D0h.
GETF-4	For any Get Feature Identifier defined in this section, the host should either clear the Namespace Identifier (NSID) field in the Get Features command to zero or set that field to FFFFFFFFh. In addition, the host may set that field to a valid NSID. If the host sends an invalid NSID value, the device shall fail the command with Invalid Field in Command.

4.15.2 DSSD Set Feature Requirements

The table below defines the scope for all DSSD specific Set Features:

Feature Identifier	Scope	Feature Name	Reference Section
C0h	NVM subsystem	Error Injection	4.15.3
C1h	NVM subsystem	Obsolete	N/A
C2h	NVM subsystem	EOL/PLP Failure Mode	4.15.5
C3h	Controller	Clear PCIe Correctable Error Counters	4.15.7
C4h	NVM subsystem	Enable IEEE1667 Silo	4.15.9
C5h	Controller	Latency Monitor	4.15.11
C6h	NVM subsystem	PLP Health Check Interval	4.15.13
C7h	NVM subsystem	DSSD Power State	4.15.15
C8h	NVM subsystem	Set Telemetry Profile	4.15.17
C9h	Controller	DSSD Asynchronous Event Configuration	4.15.19
KEY: • Namespace = The Set Feature affects a specific namespace. • Controller = The Set Feature affects the controller that is processing the command. • NVM subsystem = The Set Feature affects the NVM subsystem.			

The device shall support the following additional vendor unique Set /Get Features Identifiers.

4.15.3 Error Injection (Feature Identifier C0h) Set Feature

Feature to inject one or more error conditions to be reported by the device. If multiple Set Features commands for this feature are sent by the host, then only information from the most recent successful command is retained (i.e., subsequent commands replace information provided by previous commands).

Requirement ID	Dword	Field	Bits	Field Description
SERRI-1	0	Command Identifier (CID)	31:16	Shall be set as defined in the NVMe Specification.
SERRI-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
SERRI-3	0	Reserved	13:10	Shall be cleared to zero.
SERRI-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
SERRI-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
SERRI-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, the device shall fail the command with Invalid Field in Command.
SERRI-7	3:2	Reserved	All	Shall be cleared to zero.
SERRI-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
SERRI-9	9:6	Data Pointer (DPTR)	All	Shall contain one or two PRPs that specify a physically contiguous 4096-byte address range containing 0 to 127 Error Injection Data Structure Entries.
SERRI-10	10	Save (SV)	31	The device shall not support setting this bit to 1b. If the controller receives this Set Features command with the bit set to 1b, then the device shall abort the command with a status of Feature Identifier Not Saveable.
SERRI-11	10	Reserved	30:8	Shall be cleared to zero.
SERRI-12	10	Feature Identifier (FID)	7:0	Shall be set to C0h.
SERRI-13	11	Reserved	31:7	Shall be cleared to zero.
SERRI-14	11	Number of Error Injections	6:0	This field shall specify the number of valid Error Injection Data Entries described in the address range pointed to by the Data Pointer (DPTR) field.
SERRI-15	13:12	Reserved	All	Shall be cleared to zero.
SERRI-16	14	UUID Index	31:0	Shall be set per UUID-3.
SERRI-17	15	Reserved	31:0	Shall be cleared to zero.

Requirement ID	Description
ERRI-1	The maximum number of entries in the Number of Error Injections field shall be 127.
ERRI-2	A value of 0000000b in the Number of Error Injections field shall clear any outstanding error injection events.
ERRI-3	The error injections shall not overlap and may be listed in any order (e.g., ordering by error injection type is not required).
ERRI-4	The host shall clear any unused entries in the Error Injection data structure to zero and the device shall ignore all zeroed entries. The device shall check at least the first four bytes of each Error Injection Entry data structure to determine if it is zeroed.
ERRI-5	The device shall abort the Error Injection Set Feature command if the request contains an error injection type that is not supported or the Single Instance value for the given Error Injection Type is not valid.
ERRI-6	Once the trigger conditions specified in an Error Injection Entry are met, the device shall inject the defined error event such that the host can detect the error through either an AEN being sent, the CFS bit being set, or command being aborted.

4.15.3.1 Error Injection Entry Data Structure

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>ERRIE-1</td><td>0</td><td>Error Entry Flags</td><td>1</td><td>Error Entry Flags definition:
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:2</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>1</td><td>Single Instance: If cleared to 0b, indicates error injection is enabled until disabled. If set to 1b, indicates a single instance error injection where a single error shall be injected.</td></tr>
<tr><td>0</td><td>Error Injection Enable: If cleared to 0b, indicates error injection is disabled. If set to 1b, indicates error injection is enabled.</td></tr>
</table></td></tr>
<tr><td>ERRIE-2</td><td>1</td><td>Reserved</td><td>1</td><td>Shall be cleared to zero.</td></tr>
<tr><td>ERRIE-3</td><td>3:2</td><td>Error Injection Type</td><td>2</td><td>Error Injection type definition:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0000h</td><td>Reserved.</td></tr>
<tr><td>0001h</td><td>CPU/Controller Hang</td></tr>
<tr><td>0002h</td><td>NAND Hang</td></tr>
<tr><td>0003h</td><td>PLP Defect</td></tr>
<tr><td>0004h</td><td>Logical Firmware Error</td></tr>
<tr><td>0005h</td><td>DRAM Corruption Critical Path</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>0006h</td><td>DRAM Corruption Non-Critical Path</td></tr>
<tr><td>0007h</td><td>NAND Corruption</td></tr>
<tr><td>0008h</td><td>SRAM Corruption</td></tr>
<tr><td>0009h</td><td>HW Malfunction</td></tr>
<tr><td>000Ah</td><td>No More NAND Spares Available</td></tr>
<tr><td>000Bh</td><td>Incomplete Shutdown</td></tr>
<tr><td>000Ch - FFFFh</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>ERRIE-4</td><td>31:4</td><td>Error Injection Type Specific Definition</td><td>28</td><td>Error Injection Type specific definition.</td></tr>
</table>

4.15.3.2 Device Panic Error Injection Type

The device shall inject a device panic that the host can detect through either an AEN or the CFS bit being set. For the Device Panic type, a Single Instance value of 0 is not valid. Host shall perform the Panic Reset and Device Recovery actions specified in Error Recovery (Log Identifier C1h).

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>ERRIEDP-1</td><td>0</td><td>Error Entry Flags</td><td>1</td><td>Device Panic Error Entry Flags:<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7:2</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>1</td><td>Shall be set to 1b.</td></tr>
<tr><td>0</td><td>Shall be set to 1b.</td></tr>
</table></td></tr>
<tr><td>ERRIEDP-2</td><td>1</td><td>Reserved</td><td>1</td><td>Shall be cleared to zero.</td></tr>
<tr><td>ERRIEDP-3</td><td>3:2</td><td>Error Injection Type</td><td>2</td><td>Shall be set to the range of 0001h to 000Bh.</td></tr>
<tr><td>ERRIEDP-4</td><td>31:4</td><td>Error Injection Type Specific Definition</td><td>28</td><td>Device Panic Error Injection information:<table>
<tr><th>Byte Address</th><th>Byte Description</th></tr>
<tr><td>31:6</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>5:4</td><td>Number of Reads to Trigger Device Panic (NRTDP): Indicates the number of Read commands the device shall process and complete before triggering a device panic.</td></tr>
</table></td></tr>
</table>

4.15.4 Error Injection (Feature Identifier C0h) Get Feature

This Get Feature returns the set of error injections that are enabled on the device, GETF-4 specifies the acceptable values of the NSID field in the Get Features command. The attributes specified in Section 4.15.3.2 - Device Panic Error Injection Type are returned in Dword 0 of the completion queue entry and the Error Inject data structure specified in Section 4.15.3.1 - Error Injection Entry Data Structure is returned for each error injection in the data buffer for that command. If there are no currently enabled error injections, the data buffer returned shall contain all zeroes. The device shall clear to zero all unused entries in the Error Injection data structure.

4.15.4.1 Error Injection – Get Features Completion Queue Entry Dword 0

Requirement ID	Field	Bits	Field Description
GERRI-1	Reserved	31:7	Shall be cleared to zero.
GERRI-2	Number of Error Injections (NUME)	6:0	This field indicates the number of outstanding enabled error injections returned in the command data buffer (see Section 4.15.3.1 Error Injection Entry Data Structure for the format of the entries).

4.15.5 EOL/PLP Failure Mode (Feature Identifier C2h) Set Feature

This Set Feature defines the mode to which the device shall transition at End of Life (EOL) or on failure of the Power Loss Protection (PLP) circuitry.

Requirement ID	Description
ROWTM-1	The device shall default from the factory to Read Only Mode (ROM) (01b).

Requirement ID	Dword	Field	Bits	Field Description
SROWTM-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification version 1.4b.
SROWTM-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
SROWTM-3	0	Reserved	13:10	Shall be cleared to zero.
SROWTM-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
SROWTM-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
SROWTM-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, the device shall fail the command with Invalid Field in Command.
SROWTM-7	3:2	Reserved	All	Shall be cleared to zero.

<table>
<tr><th>Requirement ID</th><th>Dword</th><th>Field</th><th>Bits</th><th>Field Description</th></tr>
<tr><td>SROWTM-8</td><td>5:4</td><td>Metadata Pointer (MPTR)</td><td>All</td><td>Shall be cleared to zero.</td></tr>
<tr><td>SROWTM-9</td><td>9:6</td><td>Data Pointer (DPTR)</td><td>All</td><td>Shall be cleared to zero.</td></tr>
<tr><td>SROWTM-10</td><td>10</td><td>Save (SV)</td><td>31</td><td>This bit specifies that the controller shall save the End-of-Life Behavior state so that the state persists through all power states and resets. The device shall support setting this bit to 1b.</td></tr>
<tr><td>SROWTM-11</td><td>10</td><td>Reserved</td><td>30:8</td><td>Shall be cleared to zero.</td></tr>
<tr><td>SROWTM-12</td><td>10</td><td>Feature Identifier (FID)</td><td>7:0</td><td>Shall be set to C2h.</td></tr>
<tr><td>SROWTM-13</td><td>11</td><td>End of Life Behavior</td><td>31:30</td><td>Field to indicate device write behavior at End of Life (EOL) or in the event of loss of PLP functionality. See EOL-5 for definition on EOL.
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>00b</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>01b</td><td>The device shall transition to Read Only Mode (ROM) in the event of PLP failure or at EOL.</td></tr>
<tr><td>10b</td><td>The device shall transition to Write Through Mode (WTM) in the event of PLP failure and transition to Read Only Mode (ROM) at EOL.</td></tr>
<tr><td>11b</td><td>The device shall continue to operate as normal in the event of PLP failure and transition to Read Only Mode (ROM) at EOL where EOL is defined as the device no longer has enough spare blocks to support Write commands.</td></tr>
</table></td></tr>
<tr><td>SROWTM-14</td><td>11</td><td>Reserved</td><td>29:0</td><td>Shall be cleared to zero.</td></tr>
<tr><td>SROWTM-15</td><td>13:12</td><td>Reserved</td><td>All</td><td>Shall be cleared to zero.</td></tr>
<tr><td>SROWTM-16</td><td>14</td><td>UUID Index</td><td>31:0</td><td>Shall be set per UUID-3.</td></tr>
<tr><td>SROWTM-17</td><td>15</td><td>Reserved</td><td>31:0</td><td>Shall be cleared to zero.</td></tr>
</table>

4.15.6 EOL/PLP Failure Mode (Feature Identifier C2h) Get Feature

GETF-4 specifies the acceptable values of the NSID field in the Get Features command. Dword 0 of command completion queue entry.

Requirement ID	Field	Bits	Field Description
GROWTM-1	Reserved	31:3	Shall be cleared to zero.

<table>
<tr><th>Requirement ID</th><th>Field</th><th>Bits</th><th>Field Description</th></tr>
<tr><td>GROWTM -2</td><td>End of Life Behavior</td><td>2:0</td><td>Field to indicate what the device write behavior is configured for at End of Life (EOL) or in the event of loss of PLP functionality. The tables below define the required return values for each Selection (SEL) state. All other bit values are reserved.
Current state (Selection (SEL) cleared to 00b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>001b</td><td>The device will transition to Read Only Mode (ROM) in the event of PLP failure or at EOL.</td></tr>
<tr><td>010b</td><td>The device will transition to Write Through Mode (WTM) in the event of PLP failure and transition to Read Only Mode (ROM) at EOL.</td></tr>
<tr><td>011b</td><td>The device will continue to operate as normal in the event of PLP failure and transition to Read Only Mode (ROM) at EOL.</td></tr>
</table>
Default state (Selection (SEL) set to 01b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>001b</td><td>Read Only Mode (ROM) is the factory default.</td></tr>
<tr><td>010b</td><td>The Write Through Mode (WTM) is the factory default.</td></tr>
<tr><td>011b</td><td>Normal operation is the factory default for PLP failure.   Read Only Mode (ROM) is the factory default at EOL.</td></tr>
</table>
Saved state (Selection (SEL) set to 10b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>001b</td><td>The saved state is Read Only Mode (ROM).</td></tr>
<tr><td>010b</td><td>The saved state is Write Through Mode (WTM).</td></tr>
<tr><td>011b</td><td>The saved state is to operate as normal in the event of PLP failure and Read Only Mode (ROM) at EOL.</td></tr>
</table>
Capabilities (Selection (SEL) set to 11b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>101b</td><td>This feature is saveable, changeable, and not namespace specific.</td></tr>
</table>
</td></tr>
</table>

4.15.7 Clear PCIe Correctable Error Counters (Feature Identifier C3h) Set Feature

Requirement ID	Dword	Field	Bits	Field Description
CPCIE-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification version 1.4b.
CPCIE -2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
CPCIE -3	0	Reserved	13:10	Shall be cleared to zero.
CPCIE -4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
CPCIE-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
CPCIE-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
CPCIE-7	3:2	Reserved	All	Shall be cleared to zero.
CPCIE-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
CPCIE-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
CPCIE-10	10	Save (SV)	31	The device shall not support setting this bit to 1b. If the controller receives this Set Features command with the bit set to 1b, then the device shall abort the command with a status of Feature Identifier Not Saveable.
CPCIE-11	10	Reserved	30:8	Shall be cleared to zero.
CPCIE-12	10	Feature Identifier (FID)	7:0	Shall be set to C3h.
CPCIE-13	11	Clear PCIe Error Counters	31	Set to 1b to clear all PCIe correctable error counters in the SMART / Health Information Extended (Log Identifier C0h). The NVMe CLI plug-in command “clear-pcie-correctable-errors” can also perform this operation.
CPCIE-14	11	Reserved	30:0	Shall be cleared to zero.
CPCIE-15	13:12	Reserved	All	Shall be cleared to zero.
CPCIE-16	14	UUID Index	31:0	Shall be set per UUID-3.
CPCIE-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.8 Clear PCIe Correctable Error Counters (Feature Identifier C3h) Get Feature

A Get Features command for Feature Identifier C3h should not be issued because no useful information is returned. GETF-4 specifies the acceptable values of the NSID field in a Get Features command.

Requirement ID	Description
GPCIE-1	The device shall clear Dword 0 of the completion queue entry to zero.

4.15.9 Enable IEEE1667 Silo (Feature Identifier C4h) Set Feature

This Set Feature shall return an error if the OPAL Security state is not Manufactured Inactive.

Requirement ID	Dword	Field	Bits	Field Description
S1667-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVM Express Base Specification.
S1667-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
S1667-3	0	Reserved	13:10	Shall be cleared to zero.
S1667-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
S1667-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
S1667-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
S1667-7	3:2	Reserved	All	Shall be cleared to zero.
S1667-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
S1667-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
S1667-10	10	Save (SV)	31	This bit specifies that the controller shall save the IEEE1667 Silo Enable/Disable state so that the state persists through all power states and resets. The device shall support setting this bit to 1b.
S1667-11	10	Reserved	30:8	Shall be cleared to zero.
S1667-12	10	Feature Identifier (FID)	7:0	Shall be set to C4h.
S1667-13	11	Enable IEEE1667 Silo	31	If set to 0b, the IEEE 1667 silo shall be disabled no later than the next power cycle. If set to 1b, the IEEE

Requirement ID	Dword	Field	Bits	Field Description
				1667 silo shall be enabled no later than the next power cycle.
S1667-14	11	Reserved	30:0	Shall be cleared to zero.
S1667-15	13:12	Reserved	All	Shall be cleared to zero.
S1667-16	14	UUID Index	31:0	Shall be set per UUID-3.
S1667-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.10 Enable IEEE1667 Silo (Feature Identifier C4h) Get Feature

GETF-4 specifies the acceptable values of the NSID field in the Get Features command.

Dword 0 of command completion queue entry.

<table>
<tr><th>Requirement ID</th><th>Field</th><th>Bits</th><th>Field description</th></tr>
<tr><td>G1667-1</td><td>Reserved</td><td>31:3</td><td>Shall be cleared to zero.</td></tr>
<tr><td>G1667-2</td><td>IEEE1667 Silo Enabled</td><td>2:0</td><td>
<p>The tables below define the required return values for each Selection (SEL) state. All other values are illegal.</p>
<p>Current state (Selection (SEL) cleared to 00b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>000b</td><td>The IEEE1667 silo is currently disabled.</td></tr>
<tr><td>001b</td><td>The IEEE1667 silo is currently enabled.</td></tr>
</table>
<p>Default state (Selection (SEL) set to 01b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>000b</td><td>The IEEE1667 silo factory default is disabled. Unless otherwise specified, the device shall set the default value of the IEEE1667 silo to 000b (disabled).</td></tr>
<tr><td>001b</td><td>The IEEE1667 silo factory default is enabled.</td></tr>
</table>
<p>Saved state (Selection (SEL) set to 10b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>000b</td><td>The IEEE1667 silo saved state is disabled.</td></tr>
<tr><td>001b</td><td>The IEEE1667 silo saved state is enabled.</td></tr>
</table>
<p>Capabilities (Selection (SEL) set to 11b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>101b</td><td>This feature is saveable, changeable, and not namespace specific.</td></tr>
</table>
</td></tr>
</table>

4.15.11 Latency Monitor (Feature Identifier C5h) Set Feature

This configures the Latency Monitor Feature.

Requirement ID	Dword	Field	Bits	Field Description
LMSF-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification version 1.4b.
LMSF-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
LMSF-3	0	Reserved	13:10	Shall be cleared to zero.
LMSF-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
LMSF-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
LMSF-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
LMSF-7	3:2	Reserved	All	Shall be cleared to zero.
LMSF-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
LMSF-9	9:6	Data Pointer (DPTR)	All	Shall contain one PRP that specifies a physically contiguous 4096-byte address range that is 4096-byte aligned (see Section 4.15.11.1 Latency Monitoring Data Structure Entry).
LMSF-10	10	Save (SV)	31	This bit specifies that the controller shall save the data pointed to by the Data Pointer in LMSF-9 (Data Pointer (DPTR)) so that the data persists through all power states and resets. The device shall support setting this bit to 1b.
LMSF-11	10	Reserved	30:8	Shall be cleared to zero.
LMSF-12	10	Feature Identifier (FID)	7:0	Shall be set to C5h.
LMSF-13	11	Reserved	31:0	Shall be cleared to zero.
LMSF-14	12:13	Reserved	All	Shall be cleared to zero.
LMSF-15	14	UUID Index	31:0	Shall be set per UUID-3.
LMSF-16	15	Reserved	31:0	Shall be cleared to zero.

4.15.11.1 Latency Monitoring Data Structure Entry

This data structure is 4096 bytes with the following functional requirements and field format:

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td>LMDS-1</td><td>1:0</td><td>Active Bucket Timer Threshold</td><td>2</td><td>This is the value that loads the Active Bucket Timer Threshold.</td></tr>
<tr><td>LMDS-2</td><td>2</td><td>Active Threshold A</td><td>1</td><td>This is the value that loads into the Active Threshold A</td></tr>
<tr><td>LMDS-3</td><td>3</td><td>Active Threshold B</td><td>1</td><td>This is the value that loads into the Active Threshold B</td></tr>
<tr><td>LMDS-4</td><td>4</td><td>Active Threshold C</td><td>1</td><td>This is the value that loads into the Active Threshold C</td></tr>
<tr><td>LMDS-5</td><td>5</td><td>Active Threshold D</td><td>1</td><td>This is the value that loads into the Active Threshold D</td></tr>
<tr><td>LMDS-6</td><td>7:6</td><td>Active Latency Config.</td><td>2</td><td>This is the value that loads into the Active Latency Configuration.</td></tr>
<tr><td>LMDS-7</td><td>8</td><td>Active Latency Minimum Window</td><td>1</td><td>This is the value that loads into the Active Latency Minimum Window.</td></tr>
<tr><td>LMDS-8</td><td>10:9</td><td>Debug Log Trigger Enable</td><td>2</td><td>This is the value that loads into the Debug Log Trigger Enable.</td></tr>
<tr><td>LMDS-9</td><td>11</td><td>Discard Debug Log</td><td>1</td><td>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>00h</td><td>When cleared to 00h the debug log, if it exists, will not be cleared.</td></tr>
<tr><td>01h</td><td>When set to 01h the debug log will be discarded so another log can be triggered.
In addition to this, all the fields in the Set Features Data structure are valid, which will cause the Latency Monitor feature to be reset and loaded with the values from the Set Features command.
This shall discard the Debug Log regardless of whether the Latency Monitoring Feature is enabled or disabled.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Byte Address</th><th>Field</th><th># of Bytes</th><th>Field Description</th></tr>
<tr><td></td><td></td><td></td><td></td><td><table>
<tr><td>02h</td><td>When set to 02h the debug log will be discarded so another log can be triggered.
In addition to this none of the other fields of the Set Features Data structure are valid. Thus, only the debug log is discarded, and the Latency Monitor feature is not reset or loaded with any new values from the Set Features command.
This shall discard the Debug Log regardless of whether the Latency Monitoring Feature is enabled or disabled.</td></tr>
<tr><td>03h -FFh</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>LMDS-10</td><td>12</td><td>Latency Monitor Feature Enable</td><td>1</td><td>When set to 01h the Latency Monitor Feature is enabled. When cleared to 00h the Latency Monitor Feature is disabled. All other values are reserved.</td></tr>
<tr><td>LMDS-11</td><td>4095:13</td><td>Reserved</td><td>4083</td><td>Shall be cleared to zero.</td></tr>
</table>

4.15.12 Latency Monitor (Feature Identifier C5h) Get Feature

A Get Features command for Feature Identifier C5h should not be issued because no useful information is returned. GETF-4 specifies the acceptable values of the NSID field in a Get Features command.

Requirement ID	Description
LMGF-1	The device shall clear Dword 0 of the completion queue entry to zero.

4.15.13 PLP Health Check Interval (Feature Identifier C6h) Set Feature

This Set Feature defines the test interval that the device shall use to test that its Power Loss Protection (PLP) circuitry is healthy.

Requirement ID	Dword	Field	Bits	Field Description
SPLPI-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification version 1.4b.
SPLPI-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
SPLPI-3	0	Reserved	13:10	Shall be cleared to zero.
SPLPI-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.

Requirement ID	Dword	Field	Bits	Field Description
SPLPI-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
SPLPI-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
SPLPI-7	3:2	Reserved	All	Shall be cleared to zero.
SPLPI-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
SPLPI-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
SPLPI-10	10	Save (SV)	31	This bit specifies that the controller shall save the PLP Health Check Interval so that the time interval persists through all power states and resets. The device shall support setting this bit to 1b.
SPLPI-11	10	Reserved	30:8	Shall be cleared to zero.
SPLPI-12	10	Feature Identifier (FID)	7:0	Shall be set to C6h.
SPLPI-13	11	PLP Health Check Interval	31:16	The time interval between PLP health checks in minutes. If cleared to 0000h, then the health check is disabled.
SPLPI-14	11	Reserved	15:0	Shall be cleared to zero.
SPLPI-15	13:12	Reserved	All	Shall be cleared to zero.
SPLPI-16	14	UUID Index	31:0	Shall be set per UUID-3.
SPLPI-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.14 PLP Health Check Interval (Feature Identifier C6h) Get Feature

GETF-4 specifies the acceptable values of the NSID field in the Get Features command. Dword 0 of command completion queue entry.

<table>
<tr><th>Requirement ID</th><th>Field</th><th>Bits</th><th>Field description</th></tr>
<tr><td>GPLPI-1</td><td>Reserved</td><td>31:16</td><td>Shall be cleared to zero.</td></tr>
<tr><td>GPLPI-2</td><td>PLP Health Check Interval</td><td>15:0</td><td>The tables below define the required return values for each Selection (SEL) state. All other values are illegal.
Current state (Selection (SEL) cleared to 00b):<table><tr><th>Value</th><th>Description</th></tr><tr><td>0000h</td><td>The PLP Health Check is currently disabled.</td></tr></table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Field</th><th>Bits</th><th>Field description</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><td>xxxxh</td><td>The PLP Health Check Interval is currently xxxxh minutes.</td></tr>
</table>
Default state (Selection (SEL) set to 01b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0000h</td><td>The PLP Health Check Interval factory default is disabled.</td></tr>
<tr><td>000Fh</td><td>The PLP Health Check Interval factory default is 000Fh minutes.</td></tr>
</table>
Saved state (Selection (SEL) set to 10b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0000h</td><td>The PLP Health Check Interval saved state is disabled.</td></tr>
<tr><td>xxxxh</td><td>The PLP Health Check Interval saved state is xxxxh minutes.</td></tr>
</table>
Capabilities (Selection (SEL) set to 11b):
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0005h</td><td>This feature is saveable, changeable, and not namespace specific.</td></tr>
</table>
</td></tr>
</table>

4.15.15 DSSD Power State (Feature Identifier C7h) Set Feature

Requirement ID	Description
DSSDPSS-1	If the host selects an DSSD Power State via a DSSD Power State (Feature Identifier C7h) Set Feature, the device shall accept the command and run at the highest powered NVMe Power State whose Maximum Power (MP) is less than or equal to the number of the DSSD Power State in watts (see DCLP-9 and DSSDPSD-3).
DSSDPSS-2	If the host selects a DSSD Power State via DSSD Power State (Feature Identifier C7h) Set Feature and the number of that DSSD Power State is less than the Minimum DSSD Power State (see DCLP-8), the device shall abort the Set Feature command with Invalid Field in Command status.

This Set Feature causes the device to move to the given DSSD Power State:

Requirement ID	Dword	Field	Bits	Field Description
SDSSDPS-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification version 1.4b.
SDSSDPS-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.

Requirement ID	Dword	Field	Bits	Field Description
SDSSDPS-3	0	Reserved	13:10	Shall be cleared to zero.
SDSSDPS-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
SDSSDPS-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
SDSSDPS-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
SDSSDPS-7	3:2	Reserved	All	Shall be cleared to zero.
SDSSDPS-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
SDSSDPS-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
SDSSDPS-10	10	Save (SV)	31	This bit specifies that the controller shall persist the DSSD Power State through all power cycles and resets. The device shall support setting this bit to 1b.
SDSSDPS-11	10	Reserved	30:8	Shall be cleared to zero.
SDSSDPS-12	10	Feature Identifier (FID)	7:0	Shall be set to C7h.
SDSSDPS-13	11	Reserved	31:7	Shall be cleared to zero.
SDSSDPS-14	11	DSSD Power State	6:0	DSSD Power State to set in watts.
SDSSDPS-15	13:12	Reserved	All	Shall be cleared to zero.
SDSSDPS-16	14	UUID Index	31:0	Shall be set per UUID-3.
SDSSDPS-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.16 DSSD Power State (Feature Identifier C7h) Get Feature

Requirement ID	Description
DSSDPSG-1	If a DSSD Power State (Feature Identifier C7h) Get Feature command is executed and the current NVMe Power State was selected because of a DSSD Power State Set Feature command, the device shall report the DSSD Power State selected by that command (see Section 4.8.11 DSSD Power State Requirements for an example).
DSSDPSG-2	If a DSSD Power State (Feature Identifier C7h) Get Feature command is executed and the current NVMe Power State was not selected because of a DSSD Power State Set Feature command, the device shall report the lowest power DSSD Power State whose number is

Requirement ID	Description
	greater than or equal to the Maximum Power in watts of the current NVMe Power State (see Section 4.8.11 DSSD Power State Requirements for an example).

GETF-4 specifies the acceptable values of the NSID field in the Get Features command. Dword 0 of command completion queue entry.

<table>
<tr><th>Requirement ID</th><th>Field</th><th>Bits</th><th>Field Description</th></tr>
<tr><td>GDSSDPS-1</td><td>Reserved</td><td>31:7</td><td>Shall be cleared to zero.</td></tr>
<tr><td>GDSSDPS-2</td><td>DSSD Power State</td><td>6:0</td><td>The tables below define the required return values for each Selection (SEL) state. All other values are illegal.
<p>Current state (Selection (SEL) cleared to 00b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>xx</td><td>The DSSD Power State is currently xx watts as specified by DSSDPSG-1 and DSSDPSG-2.</td></tr>
</table>
<p>Default state (Selection (SEL) set to 01b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>xx</td><td>The DSSD Power State factory default is xx watts.</td></tr>
</table>
<p>Saved state (Selection (SEL) set to 10b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>xx</td><td>The DSSD Power State saved state is xx watts.</td></tr>
</table>
<p>Capabilities (Selection (SEL) set to 11b):</p>
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>0005h</td><td>This feature is saveable, changeable, and not namespace specific.</td></tr>
</table>
</td></tr>
</table>

4.15.17 Set Telemetry Profile (Feature Identifier C8h) Set Feature

4.15.17.1 Telemetry Profiles

Telemetry profiles enable the following to be configured:

- Data Area 1 Last Block
- Data Area 2 Last Block
- Statistic Area Location and Size
- Event FIFO Start and End locations and number of Event FIFOs

In the telemetry log there is the Number of Telemetry Profiles Supported field. This field enables the device to report the number of Telemetry Profiles supported. There is also a Telemetry Profile Selected field. This reports what profile is currently selected. The Telemetry Profile Selected field can be changed through Set Features. The Telemetry log can then be read to see the format of the Telemetry Profile. Having different profiles enables the device to optimize debug for different situations.

4.15.17.2 Set Telemetry Profile (Feature Identifier C8h) Set Feature

Set Feature to set the Telemetry Profile.

Requirement ID	Dword	Field	Bits	Field Description
TEL-CFG-SF-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVMe Specification.
TEL-CFG-SF-2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
TEL-CFG-SF-3	0	Reserved	13:10	Shall be cleared to zero.
TEL-CFG-SF-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
TEL-CFG-SF-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
TEL-CFG-SF-6	1	Namespace Identifier (NSID)	31:0	Shall be set to FFFFFFFFh.
TEL-CFG-SF-7	3:2	Reserved	All	Shall be cleared to zero.
TEL-CFG-SF-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
TEL-CFG-SF-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
TEL-CFG-SF-10	10	Save (SV)	31	This bit shall be set to 1b.
TEL-CFG-SF-11	10	Reserved	30:8	Shall be cleared to zero.
TEL-CFG-SF-12	10	Feature Identifier (FID)	7:0	Shall be set to C8h.
TEL-CFG-SF-13	11	Reserved	31:8	Shall be cleared to zero.
TEL-CFG-SF-14	11	Telemetry Profile Select (TPS)	7:0	The device shall collect debug data per the specified profile number.
TEL-CFG-SF-15	13:12	Reserved	All	Shall be cleared to zero.
TEL-CFG-SF-16	14	UUID Index	31:0	Shall be set per UUID-3.
TEL-CFG-SF-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.18 Set Telemetry Profile (Feature Identifier C8h) Get Feature

A Get Features command for Feature Identifier C8h should not be issued because no useful information is returned. GETF-4 specifies the acceptable values of the NSID field in a Get Features command.

Requirement ID	Description
TEL-CFG-GF-1	The device shall clear Dword 0 of the completion queue entry to zero.

4.15.19 DSSD Asynchronous Event Configuration (Feature Identifier C9h) Set Feature

This Feature controls the DSSD events that cause an asynchronous event notification to be sent to the host. If the condition for an event is true when the corresponding notice indicated in Command Dword 11 is enabled and there's an outstanding Asynchronous Event Request command, then that event is sent to the host. If the notice indicated in Command Dword 11 is disabled, then that event is not sent to the host.

Requirement ID	Dword	Field	Bits	Field Description
SDAEC-1	0	Command Identifier (CID)	31:16	Shall be set as defined in NVM Express Base Specification.
SDAEC -2	0	PRP or SGL for Data Transfer (PSDT)	15:14	Shall be cleared to 00b to specify that PRPs are used.
SDAEC -3	0	Reserved	13:10	Shall be cleared to zero.
SDAEC-4	0	Fused Operation (FUSE)	9:8	Shall be cleared to 00b.
SDAEC-5	0	Opcode (OPC)	7:0	Shall be set to 09h.
SDAEC-6	1	Namespace Identifier (NSID)	31:0	The host should either clear this to zero or set this to FFFFFFFFh. If the host sends a valid NSID value other than zero or FFFFFFFFh, then the device shall fail the command with Feature Not Namespace Specific. If the host sends an invalid NSID value, then the device shall fail the command with Invalid Field in Command.
SDAEC-7	3:2	Reserved	All	Shall be cleared to zero.
SDAEC-8	5:4	Metadata Pointer (MPTR)	All	Shall be cleared to zero.
SDAEC-9	9:6	Data Pointer (DPTR)	All	Shall be cleared to zero.
SDAEC-10	10	Save (SV)	31	This bit specifies that the controller shall save the DSSD Asynchronous Event Configuration state so that the state persists through all power states and resets. This bit shall be cleared to 0b to indicate that this feature is not saveable.
SDAEC-11	10	Reserved	30:8	Shall be cleared to zero.
SDAEC-12	10	Feature Identifier (FID)	7:0	Shall be set to C9h.
SDAEC-13	11	Reserved	31:1	Shall be cleared to zero.
SDAEC-14	11	Enable Panic Notices	0	This bit specifies whether an asynchronous event notification is sent to the host for a panic event (see EREC-5).

Requirement ID	Dword	Field	Bits	Field Description
				If this bit is set to 1b and there is an outstanding Asynchronous Event Request command, then: • The panic event is sent to the host when this condition occurs; and • The Controller Fatal Status bit in the in-band Controller Status property shall not be set to 1b when a panic event occurs; however, the Controller Fatal Status bit in the out-of-band Controller Health Data Structure shall still be set to 1b. If this bit is cleared to 0b, then: • The panic event shall not be sent to the host; and • The Controller Fatal Status bit in the in-band Controller Status property and in the out-of-band Controller Health Data Structure shall be set to 1b when a panic event occurs.
SDAEC-15	13:12	Reserved	All	Shall be cleared to zero.
SDAEC-16	14	UUID Index	31:0	Shall be set per UUID-3.
SDAEC-17	15	Reserved	31:0	Shall be cleared to zero.

4.15.20 DSSD Asynchronous Event Configuration (Feature Identifier C9h) Get Feature

GETF-4 specifies the acceptable values of the NSID field in the Get Features command. Dword 0 of command completion queue entry.

Requirement ID	Field	Bits	Field description
GDAEC-1	Reserved	31:1	Shall be cleared to zero.
GDAEC-2	Enable Controller Fatal Status Notices	0	This bit indicates whether an asynchronous event notification is sent to the host for a panic event (see EREC-5). If this bit is set to 1b and there is an outstanding Asynchronous Event Request command, then the panic event is sent to the host when this condition occurs. If this bit is cleared to 0b, then the event shall not be sent to the host. The default value of this bit following a Controller Level Reset shall be 0b.

5 PCIe Requirements

The following are PCIe requirements.

Requirement ID	Description
PCI-1	The device shall support a PCIe Maximum Payload Size (MPS) of 256 bytes or larger.
PCI-2	The device shall support modification of PCIe TLP completion timeout range as defined by the PCI Express Base Specification.
PCI-3	The vendor shall disclose the vendor-specific timeout range definition in the “Completion Timeout Ranges Supported” field in the PCI Express Base Specification.
PCI-4	The device shall support disabling PCIe Completion Timeout.
PCI-5	The device shall perform a fundamental reset when PERST# is toggled, and main power has not been removed even if firmware is hung.
PCI-6	PCIe Function Level Reset shall be supported.
PCI-7	Obsolete. See PCI-5.
PCI-8	The device shall not expose a BAR4 or a BAR5 if neither CMB nor PMR is supported.
PCI-9	The device shall support MSI-X.
PCI-10	The MSI-X table shall reside in BAR0.
PCI-11	The MSI capability shall be 4-byte aligned.
PCI-12	Any optional PCI features supported by the device not described in this document shall be clearly documented and disclosed to the customer.
PCI-13	The device shall support Completion Timeout Ranges A, B and C.
PCI-14	The device shall support Relaxed Ordering.
PCI-15	The device shall not set the Relaxed Ordering bit in transactions it initiates as a Requester that require strong ordering (i.e., transactions that would violate the NVMe protocol or result in data corruption if re-ordered shall clear the Relaxed Ordering bit to 0b). For example, the Enable Relaxed Ordering bit in transactions initiated by the device as a Requester shall be cleared to 0b in MSI/MSI-X writes, writes to the Completion Queue Entry, etc.
PCI-16	The device shall set the Relaxed Ordering bit in data DMAs, if Relaxed Ordering is enabled.
PCI-17	The device shall support all link speeds between the lowest link speed and the highest link speed, inclusive (e.g., a device compliant with the PCI Express Base Specification Revision 5.0 Version 1.0 shall support link speeds of 2.5 GT/s, 5.0 GT/s, 8 GT/S, 16 GT/s, and 32 GT/s).
PCI-18	The device shall not require Device Specific Initialization (DSI).
PCI-19	The device shall not perform a soft reset when transitioning from D3hot to D0 (i.e., No_Soft_Reset bit in the Power Management Control/Status Register shall be set to ‘1’).
PCI-20	The device shall not support BARs that map to I/O Space.
PCI-21	All BARs that map to Memory Space shall be 64-bit capable (i.e., bits 2:1 of the BAR shall be set to 10b) and shall support the full 64-bit address space.
PCI-22	The sum of the size of all BARs for a single non-SR-IOV function or SR-IOV Physical Function that maps to Memory Space shall be less than or equal to 1 MiB if neither CMB nor PMR is supported.
PCI-23	The device shall require a max of one PCI bus number (e.g., no PCIe switches are allowed).

Requirement ID	Description
PCI-24	The device’s DMA engines shall support the full 64-bit address space with uniform performance across the entire address range.
PCI-25	The Device Serial Number Extended Capability shall be supported. All functions on all ports of a device are required to return the same serial number value. The serial number shall be unique for each device and independent from all other unique identifiers on the device (e.g., for namespaces, controllers, etc.).
PCI-26	The device shall support the Advanced Error Reporting Extended Capability.
PCI-27	The device shall silently discard (following normal data link layer processing) without logging or signaling an error any PCIe VDMs targeting a function while it is in function level reset (FLR). When an FLR is initiated for a function, any already-received PCIe VDMs still being processed by that function shall be silently discarded. This requirement prohibits non-silent-discard behavior during an FLR that is allowed by the PCI Express Base spec.
PCI-28	The device shall support Address Translation Services (ATS).

5.1 Boot Requirements

Requirement ID	Description
BOOT-1	The device shall support the UEFI 2.0 or later Tianocore EDKII NVMe driver.
BOOT-2	An option ROM shall not be included.

5.2 PCIe Error Logging

The following table indicates where the PCIe physical layer error counters shall be logged.   This is in addition to the aggregated PCIe error counters defined in SMART / Health Information Extended (Log Identifier C0h).

Requirement ID	Event	Logging Mechanism
PCIERR-1	Unsupported Request Error Status (URES)	Uncorrectable Error Status Register Offset 04h in PCI Express Base Specification
PCIERR-2	ECRC Error Status (ECRCES)	
PCIERR-3	Malformed TLP Status (MTS)	
PCIERR-4	Receiver Overflow Status (ROS)	
PCIERR-5	Unexpected Completion Status (UCS)	
PCIERR-6	Completer Abort Status (CAS)	
PCIERR-7	Completion Timeout Status (CTS)	
PCIERR-8	Flow Control Protocol Error Status (FCPES)	
PCIERR-9	Poisoned TLP Status (PTS)	
PCIERR-10	Data Link Protocol Error Status (DLPES)	
PCIERR-11	Advisory Non-Fatal Error Status (ANFES)	Correctable Error Status Register Offset 10h in PCI Express Base Specification

Requirement ID	Event	Logging Mechanism
PCIERR-12	Replay Timer Timeout Status (RTS)	PCIe Correctable Error Count in the SMART / Health Information Extended (Log Identifier C0h) SMART-14.
PCIERR-13	REPLAY_NUM Rollover Status (RRS)	
PCIERR-14	Bad DLLP Status (BDS)	
PCIERR-15	Bad TLP Status (BTS)	
PCIERR-16	Receiver Error Status (RES)	

5.3 Low Power Modes

Requirement ID	Description
LPWR-1	ASPM shall be disabled after a power cycle or PCIe resets as defined by the PCI Express Base Specification.
LPWR-2	The device shall support the PCIe ASPM L1 power state.

5.4 PCIe Eye Capture

Requirement ID	Description
EYE-1	A utility shall be provided that will allow the user to decode the Eye Data field of the Physical Interface Receiver Eye Opening Measurement log page.

6 Optional Device Features

Features in this section are optional. If the feature is implemented, it shall follow the requirements in this section.

6.1 Flexible Data Placement Requirements

If Flexible Data Placement is supported, then the Generic Flexible Data Placement Requirements and a minimum of one of the Flexible Data Placement configurations shall be supported:

6.1.1 Generic Flexible Data Placement Requirements

Requirement ID	Description
FDPG-1	The device shall support the copy command and Copy Descriptor Format 0.
FDPG-2	It shall be clearly disclosed to the customer which FDP Event Types are supported and what additional fields related to this are supported with each event type (FDP Event Flags, PID, NSID, Reclaim Group Identifier, Reclaim Unit Handle Identifier, etc.)

6.1.2 FDP Single Reclaim Group Configuration

Requirement ID	Description
FDPSRG-1	The number of advertised reclaim groups (NRG) shall be 1.
FDPSRG-2	The device shall support a minimum of 8 reclaim unit handles.
FDPSRG-3	The device shall set Volatile Write Cache Present to 0b (not supported).

Requirement ID	Description
FDPSRG-4	The device shall support a configuration with a minimum of 8 reclaim unit handles of Reclaim Unit Handle Type Initially Isolated.
FDPSRG-5	The device shall support a configuration with a minimum of 8 reclaim unit handles of Reclaim Unit Handle Type Persistently Isolated.
FDPSRG-6	Single reclaim group configurations shall support a minimum of 8 namespaces.
FDPSRG-7	Single reclaim group configurations shall support multiple sector sizes in a configuration.

6.1.3 FDP Die Placement Configuration

Requirement ID	Description
FDPDP-1	The number of advertised reclaim groups (NRG) shall match the number of flash dies in the NVM subsystem.
FDPDP-2	The configuration shall support a minimum of 8 reclaim unit handles. In other words, it shall be possible to utilize (NRG) x 8 placement identifiers simultaneously.

6.2 DIX Requirements

If DIX is supported, then the following shall be supported:

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>DIX-1</td><td>The following DIX configuration shall be supported.
<table>
<tr><th>Field</th><th>Value</th></tr>
<tr><td>PIL (Protection Information Location)</td><td>0b</td></tr>
<tr><td>PI (Protection Information)</td><td>011b (Type 3)</td></tr>
<tr><td>MSET (Metadata Settings)</td><td>0b (DIX)</td></tr>
<tr><td>LBAF (LBA Format)</td><td>4K+64*</td></tr>
<tr><td>LBAFEE (LBA Format Extension Enable)</td><td>0b</td></tr>
<tr><td>PRACT (Protection Information Action)</td><td>0b and 1b</td></tr>
<tr><td>PRCHK (Protection Information Check): Guard Check (bit 2)</td><td>1b</td></tr>
<tr><td>PRCHK (Protection Information Check): Application Tag Check (bit 1)</td><td>0b</td></tr>
<tr><td>PRCHK (Protection Information Check): Reference Tag check (bit 0)</td><td>0b</td></tr>
<tr><td>STC (Storage Tag Check)</td><td>0b</td></tr>
</table>
*4 KiB logical blocks, each of which has an additional 64 bytes of metadata.</td></tr>
</table>

7 Reliability

7.1 Uncorrectable Bit Error Rate

Requirement ID	Description
UBER-1	The device shall support an Uncorrectable Bit Error Rate (UBER) of < 1 sector per 10^{17} bits read.

7.2 Power On/Off Requirements

7.2.1 Time to Ready and Shutdown Requirements

Requirement ID	Description
TTR-1	The device shall respond successfully to the Identify command within 1 second of CC.EN being set to 1b provided a Shutdown is not in progress when CC.EN is set to 1b. Some models may allow a longer time (see Section 13 Device Profiles).
TTR-2	The device shall be able to process I/O commands with a successful completion within 20 seconds of CC.EN being set to 1b. Some models require a different time (see Section 13 Device Profiles).
TTR-3	The device shall be able to process Admin commands with a successful completion as soon as CSTS.RDY = 1.
TTR-4	The device shall keep CSTS.RDY = 0 until the device is able to service NVMe Admin commands. Some models may have additional requirements before CSTS.RDY can be set to 1.(see Section 13 Device Profiles).
TTR-5	The Shutdown Notification completion (i.e., the CSTS.SHST set to 10b) shall be received within 10 seconds of a normal shutdown.
TTR-6	The device Controller shall support the CC.SHN Normal and Abrupt Shutdown Notifications.
TTR-7	When a normal shutdown is completed successfully, the device shall not enter a rebuild/recovery mode on the next power on.
TTR-8	A normal shutdown shall trigger flushing to non-volatile memory of any content within the device's internal cache(s) (e.g., SRAM/DRAM), if present, which is required to be non-volatile by this specification or any specifications required by this specification.
TTR-9	Obsolete. See NVMe-CFG-3.
TTR-10	Obsolete. Moved to PLP-1.
TTR-11	Obsolete. Moved to PLP-2.
TTR-12	Metadata rebuild due to any reason (e.g., unexpected power loss) shall not exceed 120 seconds and the device shall meet its latency requirements after this.
TTR-13	Obsolete. Moved to PLP-3.
TTR-14	Obsolete. Moved to PLP-4.
TTR-15	In case of a successful Normal shutdown operation (CC.SHN = 1 set by the NVMe device driver), no data loss is tolerated even if PLP has failed.
TTR-16	Any shutdown event shall not make the device non-functional under any conditions.

Requirement ID	Description
TTR-17	Obsolete. Moved to PLP-5.
TTR-18	Obsolete. Moved to PLP-6.
TTR-19	When the CC.SHN register is written to notify the device to shut down it shall not be assumed that power will be lost even after CC.EN is cleared to 0. Under these conditions the device shall continue to function properly based on the NVMe and PCIe Specifications.
TTR-20	The device shall be able to successfully respond to Configuration space accesses within 1s from de-assertion of PERST#.
TTR-21	All AENs defined by this specification required for supporting OCP requirements shall be enabled when setting CSTS.RDY to 1b.
TTR-22	If a PCIe Conventional Reset occurs while a device is shutdown (i.e., the CSTS.SHST field is set to 10b), then upon completion of the PCIe Conventional Reset, the device shall initialize the controller so that commands issued out-of-band that require media access are able to be processed prior to the in-band host enabling the controller. This may cause the device to transition to a state where a power loss would be treated as an unsafe shutdown. This requirement overrides the requirement in the NVM Express Base specification for it to remain safe to power off the controller until CC.EN transitions from 0b to 1b following a controller shutdown.
TTR-23	The device shall not support NVM Subsystem Shutdown.

7.2.2 Incomplete/Unsuccessful Shutdown

An incomplete/unsuccessful shutdown is a Normal or Abrupt power down that did not complete 100% of the shutdown sequence for any reason that causes the device to be unable to guarantee data/metadata integrity (e.g., firmware hang/crash, capacitor failure, PLP circuit failure, etc.).

Requirement ID	Description
INCS-1	When the power-loss protection mechanism fails for any reason while power is applied, the device shall set Critical Warning bit 4 to 1b in the SMART / Health Information (Log Identifier 02h) and transition to the write behavior as defined in the EOL/PLP Failure Mode (Feature Identifier C2h).
INCS-2	The device shall incorporate a shutdown checksum or flag as the very last piece of data written to flash to detect incomplete shutdown.
INCS-3	The checksum or flag in INCS-2 shall be used on power-up to confirm that the previous shutdown was 100% successful.
INCS-4	An incomplete shutdown shall result in an increase in the Incomplete Shutdowns field (SMART-15) in the SMART / Health Information Extended (Log Identifier C0h) and the SMART / Health Information (Log Identifier 02h) Critical Warning field shall have bit 3 set to 1b.
INCS-5	If the device is in read-only mode (i.e., bit 3 in the Critical Warning field is set to 1b in the SMART / Health Information log page), then the device shall still support NVMe Admin commands, including commands that result in a purge (see SEC-43).

Requirement ID	Description
INCS-6	If the Error Recovery (Log Identifier C1h) is supported, when the device increments the Incomplete Shutdowns field in SMART / Health Information Extended (Log Identifier C0h), it shall use the following recovery procedure at the next power up: 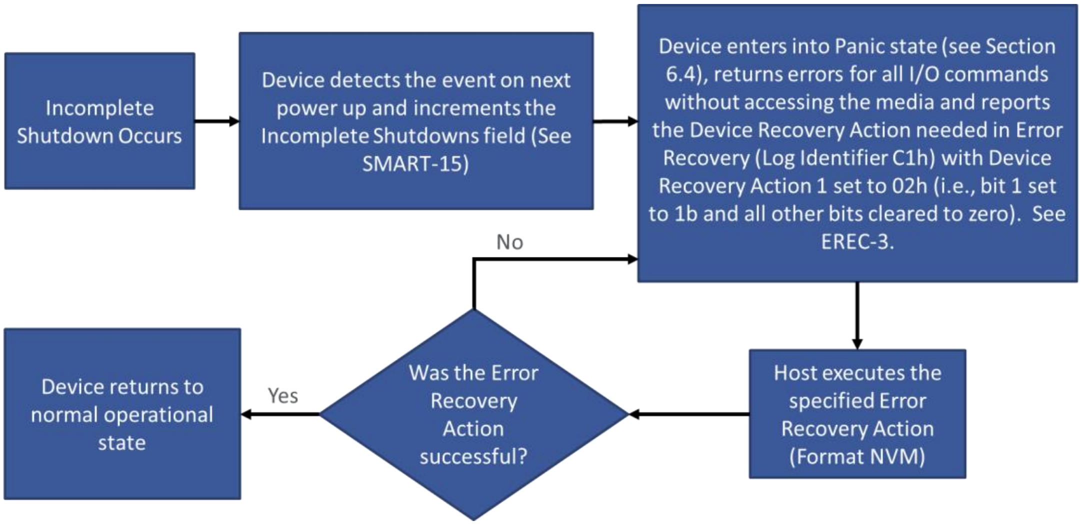 Incomplete Shutdown Occurs Device detects the event on next power up and increments the Incomplete Shutdowns field (See SMART-15) Device enters into Panic state (see Section 6.4), returns errors for all I/O commands without accessing the media and reports the Device Recovery Action needed in Error Recovery (Log Identifier C1h) with Device Recovery Action 1 set to 02h (i.e., bit 1 set to 1b and all other bits cleared to zero). See EREC-3. No Device returns to normal operational state Yes Was the Error Recovery Action successful? Host executes the specified Error Recovery Action (Format NVM)
INCS-7	If the Error Recovery (Log Identifier C1h) is not supported, when the device increments the Incomplete Shutdowns field in SMART / Health Information Extended (Log Identifier C0h), it shall use the following recovery procedure at the next power up: 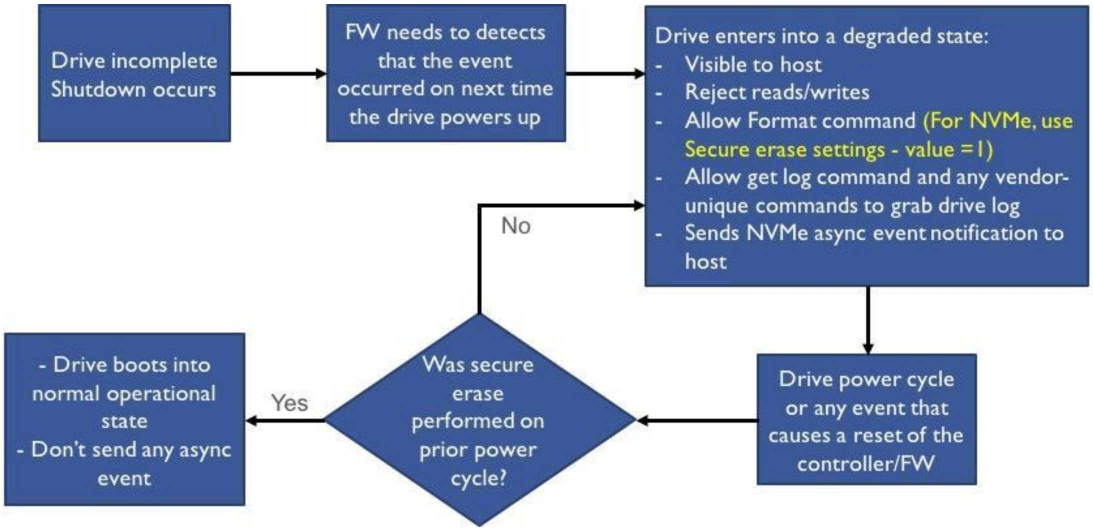 Drive incomplete Shutdown occurs FW needs to detects that the event occurred on next time the drive powers up Drive enters into a degraded state: - Visible to host - Reject reads/writes - Allow Format command (For NVMe, use Secure erase settings - value = 1) - Allow get log command and any vendor-unique commands to grab drive log - Sends NVMe async event notification to host No - Drive boots into normal operational state - Don't send any async event Yes Was secure erase performed on prior power cycle? Drive power cycle or any event that causes a reset of the controller/FW

7.3 End to End Internal Data Protection

Requirement ID	Description
E2E-1	All user data and metadata shall be protected using overlapping protection mechanisms throughout the entire read and write path in the device including all storage elements (registers, caches, SRAM, DRAM, NAND, etc.).

Requirement ID	Description
E2E-2	At least one bit of correction and 2 bits of detection is required for all memories. This shall be for all memories regardless of function.
E2E-3	The entire DRAM addressable space shall be protected with at least one-bit correction and 2 bits of detection scheme (SECDED). This includes but is not limited to the following: • Flash translation layer (FTL). • Mapping tables (including metadata related to deallocated LBAs). • Journal entries. • Firmware scratch pad. • System variables. • Firmware code.
E2E-4	Silent data corruption shall not be tolerated under any circumstances.
E2E-5	The device shall include a mechanism to protect against returning the data from the wrong logical block address (LBA), including previous copies from the same LBA, to the host. It is acceptable if the device stores additional/modified information to provide protection against returning wrong data to the host. Device shall perform host LBA integrity checking on all transfers to and from the media.
E2E-6	All device metadata, firmware, firmware variables, and other device system data shall be protected by at least a single bit detection scheme.
E2E-7	Any memory buffers that are utilized to accelerate data transfer (read-ahead buffers for example) shall follow the protection scheme outlined in E2E-5.

7.4 Behavior on Firmware Crash, Panic or Assert

Requirement ID	Description
CRASH-1	While the device is in a failed state in which it can no longer accept Write commands (e.g., firmware crash, panic or assert), the device shall allow read access only if it can guarantee data integrity.
CRASH-2	While the device is in a failed state, that is not a controller failure (e.g., firmware crash, panic or assert), the device shall still support NVMe Admin commands including the ability to read any failure logs from the device to determine the nature of the failure. The device may fail Admin commands sent by the host prior to completion of the Error Recovery workflow.
CRASH-3	Upon device completion of the actions specified in Panic Reset Action (see EREC-2) and Device Recovery Action (see EREC-3), that successfully recovers the device from a failed state and shall function normally as specified in this specification including read and write access at full performance.
CRASH-4	Obsolete
CRASH-5	The device shall have a mechanism (e.g., watchdog timer) to allow recovery from hangs of the firmware.

7.5 Annual Failure Rate (AFR)

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>REL-1</td><td>The device shall meet an MTBF of 2.5 million hours (AFR of <= 0.35% per JEDEC JESD 218) under the following environmental conditions:
<table>
<tr><th>Specification</th><th>Environment</th><th>Requirement</th></tr>
<tr><td rowspan="2"><b>Temperature</b></td><td>Operational</td><td>● 0˚C to 50˚C (32˚F to 122˚F)
● < 20°C (68˚F) per hour gradient</td></tr>
<tr><td>Non-Operational</td><td>● -40˚C to 70˚C (-40˚F to 158˚F)
● < 30˚C (86˚F) per hour gradient</td></tr>
<tr><td rowspan="2"><b>Humidity</b></td><td>Operational</td><td>● 10% to 90% non-condensing
● Yearly weighted average: < 80% RH
○ 90% of year: < 80%
○ 10% of year: 80% to 90%
● Maximum dewpoint: 29.4°C (85°F)</td></tr>
<tr><td>Non-Operational</td><td>● 5% to 95% non-condensing
● 38°C (100.4°F) maximum wet bulb temperature</td></tr>
</table></td></tr>
<tr><td>REL-2</td><td>The device shall meet an MTBF of 2.0 million hours (AFR of <= 0.44% per JEDEC JESD 218) under the following environmental conditions:
<table>
<tr><th>Specification</th><th>Environment</th><th>Requirement</th></tr>
<tr><td rowspan="2"><b>Temperature</b></td><td>Operational</td><td>● 0˚C to 55˚C (32˚F to 131˚F)
● < 20°C (68˚F) per hour gradient</td></tr>
<tr><td>Non-operational</td><td>● -40˚C to 70˚C (-40˚F to 158˚F)
● < 30˚C (86˚F) per hour gradient</td></tr>
<tr><td rowspan="2"><b>Humidity</b></td><td>Operational</td><td>● 10% to 90% non-condensing
● Maximum dewpoint: 29.4°C (85°F)</td></tr>
<tr><td>Non-operational</td><td>● 5% to 95% non-condensing
● 38°C (100.4°F) maximum wet bulb temperature</td></tr>
</table></td></tr>
<tr><td>REL-3</td><td>Supplier shall provide the temperature conditions used to determine the MTBF.</td></tr>
<tr><td>REL-4</td><td>Supplier shall provide AFR de-rating curves for the Temperature range shown in requirement REL-2 for up to 70˚C (158˚F).</td></tr>
<tr><td>REL-5</td><td>The AFR targets in REL-1 (MTBF of 2.5 million hours) and REL-2 (MTBF of 2.0 million hours) shall be maintained up to a continuous reported composite temperature of 77˚C (170 °F) (WCTEMP) with less than 1% of the device power on time above WCTEMP.</td></tr>
<tr><td>REL-6</td><td>The device shall have no exposed raw copper on any component. All copper surfaces, pads, fingers, etc. on all components shall be covered with a minimally reactive coating (e.g., resin packaging material on ICs, gold plating on PCB connectors/pads, silver-tin coatings on component leads and IC balls).</td></tr>
</table>

7.6 Background Data Refresh

Requirement ID	Description
BKGND-1	The device shall support background data refresh while the device is powered on to ensure there is no data-loss due to power-on retention issues.
BKGND-2	The device shall be designed and tested to support the normal NAND operating temperature.
BKGND-3	Background data refresh shall cover the entire device and be designed to continuously run in the background and not just during idle periods.

7.7 Wear-leveling

Requirement ID	Description
WRL-1	The device shall utilize the entire Endurance Group media capacity range whenever the device needs to wear-level a block. The device shall not restrict the wear-leveling range to a subset of the Endurance Group media capacity unless otherwise specified. If the device does not support Endurance Groups, it shall wear-level across the entire physical media.

7.8 Power Loss Protection

Requirement ID	Description
PLP-1	The device shall support full power-loss protection for all acknowledged data and metadata.
PLP-2	The Power-loss protection health check shall not impact I/O latency and performance.
PLP-3	Power-loss protection health check shall be performed by the device at power-on (prior to accepting any writes) and at least once every time interval as specified in SPLPI-13.
PLP-4	While performing the power-loss protection health check, the device shall still have enough energy to be able to handle an abrupt power loss properly.
PLP-5	The firmware/hardware algorithm shall deploy safeguards to prevent t false detection of power loss protection failure. Example of a false detection would be a glitch in any of the power loss circuitry readings which would cause a transient event to trigger a false power loss protection failure when the power loss protection hardware is healthy.
PLP-6	The device shall implement a power-loss protection (PLP) health check which can detect the capacitor holdup energy margin reported in the Capacitor Health field specified in SMART-19 in SMART / Heath Information Extended (Log Identifier C0h). The PLP health check shall not just check for open/short capacitor conditions but shall measure the true available margin energy.
PLP-7	The factory default PLP Health Check Interval shall be 15 minutes (see GPLPI-2).
PLP-8	The default PLP Health Check Interval of 15 minutes (see GPLPI-2) shall not cause the device to violate REL-1 or REL-2.
PLP-9	The device shall use the following workflow for the PLP Health Check (see INCS-1) for PLP AEN details):

Requirement ID	Description
	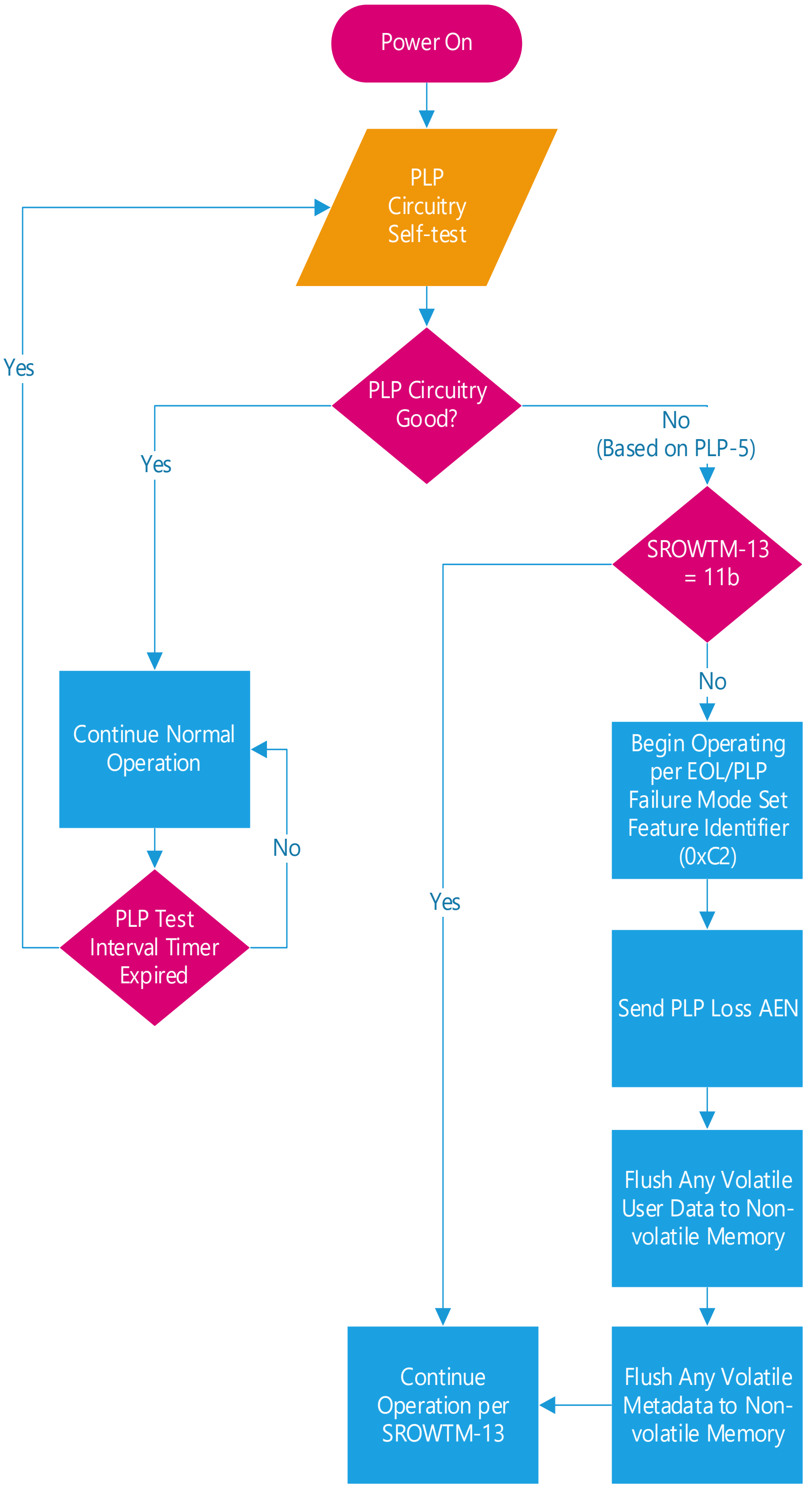 Power On PLP Circuitry Self-test PLP Circuitry Good? Yes No (Based on PLP-5) SROWTM-13 = 11b No Yes Continue Normal Operation PLP Test Interval Timer Expired Yes No Begin Operating per EOL/PLP Failure Mode Set Feature Identifier (0xC2) Send PLP Loss AEN Flush Any Volatile User Data to Non-volatile Memory Flush Any Volatile Metadata to Non-volatile Memory Continue Operation per SROWTM-13

7.9 Device Limits

Requirement ID	Description
DEVLMT-1	The device shall not have any architectural restrictions, other than P/E cycle restrictions of the NAND, on the number of times any of the following events can occur: • Firmware downloads and activation supported (see FWUP-2); or • Changing password when taking/changing ownership via TCG; or • Commands that result in a purge (see SEC-43); or • Power cycles; or • Set Features command/Get Features command (including power state changes); or • Log page or debug log retrievals. This is not meant to be a test until fail requirement.

8 Endurance

8.1 Endurance Data

Requirement ID	Description
ENDUD-1	The device documentation shall include the number of physical bytes able to be written to the device assuming a write amplification of 1. The units should be gigabytes (10^9 bytes).
ENDUD-2	The preconditioning steps to test device performance at end-of-life are: • 50/50 Read/Write workload (by number of I/Os). • 4kiB Read accesses aligned to 4kiB boundaries. • 128kiB Write accesses aligned to 128kiB boundaries. • Random pattern of Read addresses. • Sequential pattern of Write addresses. • 100% active range. • 80% full device (80% data, 20% free space). • 0% compressible data. • Ambient temperature 35°C (95°F). • Short-stroked device if capacity is 2TB or greater (see EOL-2).
ENDUD-3	The Percentage Used in the SMART / Heath Information (Log Identifier 02h) shall track linearly with bytes written and at 100% it shall match the EOL value specified in ENDUD-1.

8.2 Retention Conditions

Since there are several factors that impact the device endurance, the table below provides the requirements for the datacenter environment.

Requirement ID	Description
RETC-1	Non-Operational (Powered-off) device data retention time (end of life) shall be at least 1 month at 40°C (104°F). See Section 13 Device Profiles for specific retention requirements. Specific Devices Profiles may have longer data retention time requirements.

Requirement ID	Description
RETC-2	Operating (Powered-on) data retention shall be at least 7 years. For purposes of this requirement, the assumption is that the Terabytes Written (TBW) capability of the devices is used linearly over the lifetime. This requirement does not imply any specific warranty period.
RETC-3	The device shall not throttle its performance based on the endurance metric (endurance throttling).

8.3 Shelf Life

Requirement ID	Description
SLIFE-1	A new device may be kept as a datacenter spare and therefore shall be fully functional even if it sits on the shelf for at least 5 years at 25°C (77°F) before getting installed in the server. The device shall be new in box factory state.
SLIFE-2	Shelf-life shall be documented and provided to the customer.

8.4 End-of-Life (EOL)

Requirement ID	Description
EOL-1	Several types of samples are required for EOL testing: • Beginning of Life (Short stroked if required by EOL-2). • End of Life (Short stroked if required by EOL-2). • End of Life (Not short stroked if different than #2). EOL is defined as the Total Bytes Written (TBW) specification has been surpassed (see ENDUD-1) or the non-volatile media endurance limit (e.g., NAND cycling limit) has been reached; whichever is earliest.
EOL-2	On 2 TB or larger devices, there shall be a method to "short stroke" the device. Media reserved for background operations shall be proportionally adjusted. Short Stroke firmware capacity shall be reduced to a range of 10%-3% of full capacity. Short stroke shall cover all the channels/dies/planes while maintaining the native performance of the device.
EOL-3	Obsolete.
EOL-4	The device shall continue to operate in a read/write mode until the conditions in EOL-5 are reached.
EOL-5	The device shall switch to Read Only Mode (ROM) when the Available Spare field in the SMART / Health Information log page (Log Identifier 02h) reaches 0%. A value of 0% represents the device state where there is an insufficient number of spare blocks to support Host writes. The Available Spare field shall report a value of 0% before the device exhausts all available spare blocks so that there are still sufficient available spare blocks remaining to handle blocks going bad during read operations. When the device switches to ROM, bit 2 and bit 3 of the Critical Warning field in the SMART / Health Information log page shall be set to 1b. If the value reported in the Available Spares field falls below the value reported in the Available Spare Threshold field, then the device shall set bit 0 of the Critical Warning field in

Requirement ID	Description
	the SMART / Health Information log page to 1b and generate a SMART / Health Status event if that event is enabled and an Asynchronous Event Request command is outstanding.
EOL-6	The device shall have enough spare blocks so that the Percentage Used field value in the SMART / Health Information log page reaches 100% prior to the Available Spare field value falling below the Available Spare Threshold field value.
EOL-7	The update granularity for the Percentage Used field and Available Spare field in the SMART / Health Information log page shall be 1%.

9 Thermal

9.1 Data Center Altitude

Requirement ID	Description
THERM-1	Support for devices in data centers located at an altitude of up to 10,000 feet above sea level is required (e.g., thermals, cosmic rays, etc.). When considering thermals, there shall be no de-rating up to 6,000 feet above sea level. Above 6,000 feet the derating shall be 0.9°C (1.6°F)/1000ft.
THERM-2	A thermal study with each platform is required and shall be provided to the customer. The thermal design shall be validated up to 35°C (95°F) ambient temperature for the platform with a worst-case airflow of 1.5 meters per second at sea level.
THERM-3	The device shall operate normally with relative humidity to be between 10% and 90%.

9.2 Thermal Throttling

Requirement ID	Description
TTHROTTLE-1	The device shall implement a thermal throttling mechanism to protect the device in case of a failure or an excursion that causes the device's temperature to increase above its maximum specified temperature. This mechanism shall check device temperature at a frequency sufficient to ensure that the thermal throttling mechanism protects the device. This temperature check frequency shall be at least once per second.
TTHROTTLE-2	When the device begins throttling performance due to an over temperature condition, an Asynchronous Event Request shall be completed with the Asynchronous Event Type field set to 001b (SMART / Health status) and the Asynchronous Event Information field set to 01h (Temperature Threshold), if that event is enabled and an Asynchronous Event Request command is outstanding. The device shall set bit 1 of the SMART / Health Information (Log Identifier 02h) Critical Warning field to 1b.
TTHROTTLE-3	Thermal throttling shall only be engaged under conditions that are able to cause significant reduction in device service life, such as excessive server ambient temperature. The required behavior is illustrated below:

Requirement ID	Description
	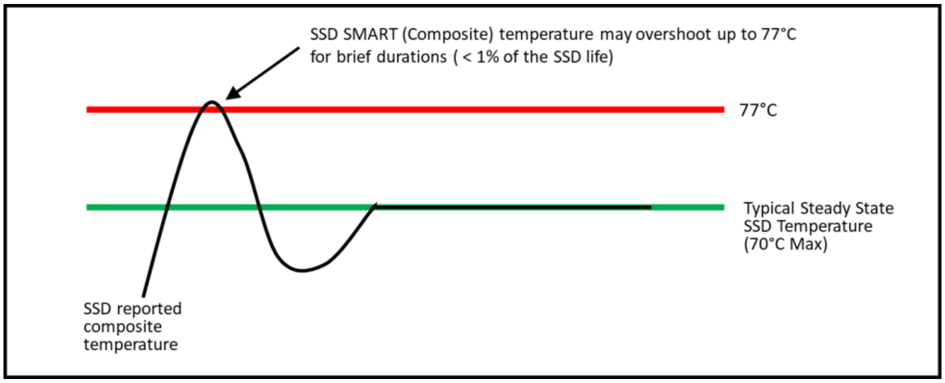
TTHROTTLE-4	The firmware algorithm shall deploy safeguards to prevent a false activation of either thermal throttling or thermal shutdown. An example of a false activation would be a glitch in any of the sensor readings which would cause the composite temperature to reach the thermal throttling or thermal shutdown limit.
TTHROTTLE-5	A composite temperature of 77°C (170.6°F) shall be used to start throttling. Single stage and multi-stage throttling are both allowed.
TTHROTTLE-6	Thermal throttling shall not start based on the rate of temperature increase or slew rate.
TTHROTTLE-7	When the device is in the thermal throttling state and the temperature drops back down below 75°C (167°F), the device shall exit the thermal throttling state and an Asynchronous Event Request shall be completed with the Asynchronous Event Type field set to 001b (SMART / Health status) and the Asynchronous Event Information field set to 01h (Temperature Threshold), if that event is enabled and an Asynchronous Event Request command is outstanding. The device shall clear bit 1 of the SMART / Health Information log page (Log Identifier 02h) Critical Warning field to 0b.
TTHROTTLE-8	When the device reaches a critical temperature on any component it shall report a composite temperature of 85°C (185°F) and an Asynchronous Event Request shall be completed with the Asynchronous Event Type field set to 001b (SMART / Health status) and the Asynchronous Event Information field set to 01h (Temperature Threshold), if that event is enabled and an Asynchronous Event Request command is outstanding, before the device shuts off to protect itself.
TTHROTTLE-9	The device shall report a value of 015Eh (77°C) in the Warning Composite Temperature Threshold (WCTEMP) field of the Identify Controller Data structure.
TTHROTTLE-10	The device shall report a value of 0166h (85°C) in the Critical Composite Temperature Threshold (CCTEMP) field of the Identify Controller Data structure.
TTHROTTLE-11	If the Composite Temperature is greater than or equal to 77°C (170.6°F), the device shall update the Warning Composite Temperature Time field of the SMART / Health Information (Log Identifier 02h) once per second until one of the following occurs: • The device Composite Temperature falls below 75°C (167°F); or • The device shuts down due to an over Critical Temperature condition; or • The device is power cycled.

Requirement ID	Description
TTHROTTLE-12	The temperature limits of internal components within the device shall govern the thermal throttling algorithms. In addition, the thermal throttling algorithms shall prevent excessive case temperature (see TTHROTTLE-14).
TTHROTTLE-13	In all throttling states including the lowest throttle state, the device shall not enter a read-only mode and shall continue to attempt to process all commands, including commands that write media, unless the device shuts down, e.g., due to a Critical Temperature condition.
TTHROTTLE-14	The critical device case (touch) temperature shall be the TS2 limits as specified in IEC 62368-1 3rd Edition or later. The TS2 100°C limit is prohibited for some form factors (see FFM2-15, FFE1S-13, and FFE1L-13). The TS2 80°C limit applies to all form factors, including those for which the TS2 100°C limit is prohibited.

9.3 Temperature Reporting

Requirement ID	Description
TRPT-1	The device shall expose the current raw sensor readings from all the sensors on the device.
TRPT-2	The device's device-to-device composite temperature variation shall be no more than +/- 1 degrees C. Two different devices shall not report a composite temperature greater than 2 degrees apart under the same environmental conditions, slot location, and workload.
TRPT-3	A single device's composite temperature shall not vary by more than +/-1°C (1.8°F) degrees once it is in a steady state under the same environmental conditions, slot location, and workload.
TRPT-4	The supplier shall provide to the customer the equation, settings, and thresholds used to calculate the composite temperature.
TRPT-5	When calculating the composite temperature, the device shall take into account the temperature of all of the major thermal components (SOC, DRAM, NAND, PMIC, etc.).
TRPT-6	The reported composite temperature shall be the same for all controllers in the NVM subsystem.

9.4 Thermal Shutdown

Requirement ID	Description
THRMS-1	The device shall shut down to protect itself against data loss or damage due to extreme temperatures. The shutdown temperature value shall be at a composite temperature of 85°C (185°F). Shutdown shall preserve all user data.

10 Form Factor Requirements

10.1 Generic Form Factor Requirements

Requirement ID	Description
GFF-1	The device shall be compliant to PCI Express Base Specification, Revision 4.0, Version 1.0 (or later).
GFF-2	The vendor shall provide a PCIe compliance report that demonstrates compliance.
GFF-3	The device shall support lane reversal with all lanes connected or partially connected lanes on both the lower-numbered lanes and higher-numbered lanes of the device (e.g., a x4 device shall support lane reversal for x4, x2, and x1 connections on both the lower-numbered lanes and higher-numbered lanes of the device).
GFF-4	The device shall train to x1 when the upstream device provides only one lane, to x2 when 2 lanes are provided and to x4 when 4 lanes are provided. This requirement applies to both the lower-numbered lanes and higher-numbered lanes of the device.
GFF-5	The device shall support hot swap on form factors that support hot swap.
GFF-6	A CAD file of each supported form factor shall be provided to the customer.

10.2 Power Consumption Measurement Methodology

Requirement ID	Description
PCM-1	The device Max Average Power (MAP) consumption for any workload shall not exceed the Maximum Power (MP) of the current NVMe Power State over a 1s window with a sampling rate of 2ms or better. The measurement duration shall be at least 15 minutes on a pre-conditioned device. This requirement defines the measurement methodology for Maximum Sustained Device Power as defined in SFF TA-1009 revision 3.0.
PCM-2	The device peak power for any workload shall be measured over a 100uS window with a sampling interval of 4uS or smaller. The measurement duration shall be at least 15 minutes on a pre-conditioned device.
PCM-3	For devices, whose max average power consumption is less than or equal to 25W, the peak power shall be no more than 1.5x the max average power. For devices, whose max average power consumption is greater than 25W but less than or equal to 29W the peak power shall be no more than 37.5W. For devices, whose max average power consumption is greater than 29W, the peak power shall be no more than 1.3x the max average power. The max average power is defined in PCM-1.

10.3 Power Levels

Requirement ID	Description
PWR-1	The Power Management (Feature Identifier 02h) shall be supported and the NVMe Power State Descriptor table in Identify Controller Data Structure bytes 3071:2048 shall be filled out per the NVMe 1.4b specification.
PWR-2	Obsolete. See DCLP-9 and Section 4.8.11.1 DSSD Power State Descriptor.

Requirement ID	Description
PWR-3	The method of measurement for Maximum Average Power (MAP) is defined in PCM-1. MAP values are reported in the Maximum Power (MP) field of NVMe Power State Descriptors.
PWR-4	Power state entries above the maximum rated power envelope of the device shall not be populated in the NVMe Power State Descriptor table in Identify Controller Data Structure bytes 3071:2048.
PWR-5	The Set Features for Power Management (Feature Identifier 02h) with the SV bit 31 in Command Dword 10 shall be supported so that the power level can be set and will be saved across power cycles.
PWR-6	The device, regardless of form factor or capacity, shall have an idle power of 5 Watts or less per European regulation.
PWR-7	The device shall not consume more power than the PCI-SIG Slot Capabilities register Slot Power Limit regardless of other settings (e.g., DSSD Power State, NVMe Power State).
PWR-8	The device shall not use the same Maximum Power (MP) value in more than one supported NVMe Power State Descriptor unless otherwise specified.

10.4 M.2 Form Factor Requirements

Requirement ID	Description
FFM2-1	The device shall adhere to the M.2 specification with a size of 22mm x 110mm.
FFM2-2	The bottom-side height shall not exceed 1.5mm.
FFM2-3	The top-side height shall not exceed 3.2mm. Some models may need to be thinner (see Section 13 Device Profiles).
FFM2-4	The device shall use an M key.
FFM2-5	The device shall support a minimum of PCIe Gen4 x4.
FFM2-6	The device shall support driving an activity LED through the connector via LED_1#.
FFM2-7	The LED should be lit solidly when power is applied and flashing when there is traffic going to the SSD.
FFM2-8	The device shall not use any pins that are defined in the m.2 specification for vendor unique functionality.
FFM2-9	The device shall support a protection scheme that protects against NAND block level failures.
FFM2-10	The protection scheme must also support NAND plane level failures without data or metadata loss.
FFM2-11	The Label shall be placed on the top side of the device.
FFM2-12	The device electricals shall follow the SMBus connection as described below and in PCI SIG M.2 3.0 1.2 or later.
FFM2-13	The device’s SMBus protocol shall comply with version 3.1 (see *http://smbus.org/specs/SMBus_3_1_20180319.pdf*).
FFM2-14	The default max average power for M.2 devices shall not exceed 8.25W and the peak power shall comply to PCM-3.

Requirement ID	Description
FFM2-15	The case (touch) temperature shall not exceed 80°C. The 100°C TS2 limit does not apply (see TTHROTTLE-14).

10.5 E1.S Form Factor Requirements

Requirement ID	Description
FFE1S-1	The device shall adhere to the latest revision of SFF-TA-1006.
FFE1S-2	At a minimum, the device shall support PCIe Gen4 x4.
FFE1S-3	The device shall support activity and error LEDs.
FFE1S-4	The activity LED shall be lit solidly when power is applied and flashing when there is traffic going to the device.
FFE1S-5	The device shall support a protection scheme that protects against NAND block level failures.
FFE1S-6	The protection scheme must also support NAND plane level failures without data or metadata loss.
FFE1S-7	The amber LED shall meet the requirements specified in SFF TA-1009. The functioning of the Amber LED shall be independent of the 12V, 3.3Vaux and the state of the PWRDIS pin.
FFE1S-8	The thermal performance of the 9.5mm, 15mm, and 25mm cases and their associated pressure drops shall be provided.
FFE1S-9	The PWRDIS pin shall be supported.
FFE1S-10	The SMBus electrical connections shall follow the "DC Specification For 3.3V Logic Signaling" as defined in SFF-TA-1009 revision 2.0. Including Vih1 with a max of 3.465V.
FFE1S-11	The device's SMBus protocol shall comply to version 3.1 (see *http://smbus.org/specs/SMBus_3_1_20180319.pdf*).
FFE1S-12	A x4 device shall only have a 1C connector.
FFE1S-13	The case (touch) temperature shall not exceed 80°C. The 100°C TS2 limit does not apply (see TTHROTTLE-14).

10.6 E1.L Form Factor Requirements

Requirement ID	Description
FFE1L-1	The device shall adhere to the latest revision of SFF-TA-1007.
FFE1L-2	At a minimum, the device shall support PCIe Gen4 x4.
FFE1L-3	The device shall support activity and error LEDs.
FFE1L-4	The activity LED shall be lit solidly when power is applied and flashing when there is traffic going to the device.
FFE1L-5	The device shall support a protection scheme that protects against NAND block level failures.
FFE1L-6	The protection scheme must also support NAND plane level failures without data or metadata loss.

Requirement ID	Description
FFE1L-7	The amber LED shall meet the requirements specified in SFF TA-1009. The functioning of the Amber LED shall be independent of the 12V, 3.3Vaux and the state of the PWRDIS pin.
FFE1L-8	The thermal performance of the 9.5mm and 18mm cases and their associated pressure drops shall be provided.
FFE1L-9	The PWRDIS pin shall be supported.
FFE1L-10	The SMBus electrical connections shall follow the "DC Specification For 3.3V Logic Signaling" as defined in SFF-TA-1009 revision 2.0. Including Vih1 with a max of 3.465V.
FFE1L-11	The device's SMBus protocol shall comply to version 3.1 (see *http://smbus.org/specs/SMBus_3_1_20180319.pdf*).
FFE1L-12	A x4 device shall only have a 1C connector.
FFE1L-13	The case (touch) temperature shall not exceed 80°C. The 100°C TS2 limit does not apply (see TTHROTTLE-14).

10.7 E3 Form Factor Requirements

Requirement ID	Description
FFE3-1	The device shall adhere to the latest revision of SFF-TA-1008.
FFE3-2	The device shall adhere to the latest revision of SFF-TA-1009.
FFE3-3	The device shall support Separate Reference Clocks with Independent SSC (SRIS) requirements of the PCI Express Base Specification. Devices shall support SRIS detection as described in SFF-TA-1009. The SRIS support shall include a clock tolerance of a 5600-ppm difference for separate reference clocks. The device shall support switching between SRIS and common clock on assertion/de-assertion of PERST# with or without a power cycle as described in SFF-TA-1009.
FFE3-4	The device shall be compliant to PCI Express Base Specification, Revision 5.0, Version 1.0 (or later).

10.8 SFF-8639 (U.2/U.3) Form Factor Requirements

Requirement ID	Description
FF8639-1	The device shall adhere to SFF-TA-1001 revision 1.1 or later and the PCI Express SFF-8639 Module Specification, Revision 4.0, Version 1.0. Compliance to any later Revision of the PCI Express SFF-8639 Module Specification is prohibited.
FF8639-2	The device shall support Separate Reference Clocks with Independent SSC (SRIS) requirements of the PCI Express Base Specification. Devices shall support SRIS detection as described in the PCI Express SFF-8639 Module Specification. The SRIS support shall include a clock tolerance of a 5600-ppm difference for separate reference clocks.

Requirement ID	Description
FF8639-3	All 2.5 SFF NVMe devices shall support the activity LED function on Pin 11 of the SFF-8639 connector. Pin 11 shall assert/de-assert (see FF8639-4) while processing one or more commands. Once asserted due to command processing, pin 11 shall remain asserted for 50ms and then shall be de-asserted for 50ms. This assert/de-assert behavior shall continue while command processing continues (i.e., The LED is usually off, but flashes on and off while commands are processed). Once pin 11 is de-asserted for 50ms, it shall remain de-asserted as long as there is no command processing activity.
FF8639-4	Active State (pin asserted and LED on) is defined as logic "0" (i.e., < 0.3 volts) on Pin 11. Inactive State is defined as logic "1" (i.e., > 3.3 volts) on Pin 11. A device's internally initiated background activity or maintenance routine that is not commanded to be performed from the host shall not cause the LED to be in the Active State (logic "0" level).

11 Management Support

11.1 NVMe Basic Management Command (Appendix A) Requirements

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>SMBUS-1</td><td>The device shall support the NVMe Basic Management Command as defined in Appendix A of the NVMe Management Interface 1.2c specification: (<em>https://nvmexpress.org/wp-content/uploads/NVM-Express-Management-Interface-Specification-1.2c-2022.10.06-Ratified.pdf</em>).</td></tr>
<tr><td>SMBUS-2</td><td>SMBus Block Read protocol and Byte Read protocol shall be supported. As an override of the NVMe Management Interface specification, SMBus Block Write protocol and Byte Write protocol shall also be supported.</td></tr>
<tr><td>SMBUS-3</td><td>The device shall implement the SMBus format as show in Section 11.2 - NVMe Basic Management Command (Appendix A) Data Format.</td></tr>
<tr><td>SMBUS-4</td><td>Unless otherwise noted, the default value for the Firmware Update Flags field (byte 91) in the SMBus Data structure shall be set to FFh.</td></tr>
<tr><td>SMBUS-5</td><td>The Secure Boot Failure Feature Reporting Supported bit at offset 243 shall be supported and set to 1b.</td></tr>
<tr><td>SMBUS-6</td><td>When there is a secure boot failure the device shall report the failure with the following behavior:
<table>
<tr><th>Bit</th><th>Description</th></tr>
<tr><td>7</td><td>Shall be set to 1b. See Command Code 242 for the definition of Secure Boot Failure Reporting.</td></tr>
<tr><td>6</td><td>Shall be set to 1b. See Command Code 242 for the definition of Secure Boot Failure Reporting.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td><table>
<tr><td>5</td><td>Shall be set to 1b if OCP Recovery/ Platform Root-of-Trust for Recovery codes are supported. If this bit is set to 1b then a valid Recovery code shall be entered in byte 244. These codes are defined in the <i>Platform Root-of-Trust for Recovery</i>. Use of these codes is explained in the <i>OCP Hardware Secure Boot White Paper</i>.</td></tr>
<tr><td>4:0</td><td>Reserved. Shall be cleared to zero.</td></tr>
</table></td></tr>
<tr><td>SMBUS-7</td><td>The device shall take no longer than the CAP.TO timeout value to produce stable SMBus output through the NVMe Basic Management Command (if supported).</td></tr>
<tr><td>SMBUS-8</td><td>If the device has set a non-zero value in the Panic ID field (Command Code 154, byte 01h), then the device shall clear the Drive Functional field to 0b (Command Code 0, byte 01h, Bit 5b).</td></tr>
<tr><td>SMBUS-9</td><td>The device shall generate the PEC values specified for each command code in the SMBus data format described in Section 11.2 - NVMe Basic Management Command (Appendix A) Data Format when the host issues a Block Read.</td></tr>
<tr><td>SMBUS-10</td><td>The device shall check the PEC value sent when the host issues a Block Write and only process the message if the PEC value matches the SMBus data format specified value described in Section 11.2 - NVMe Basic Management Command (Appendix A) Data Format. The device is encouraged to issue a NACK if the PEC value is not correct.</td></tr>
<tr><td>SMBUS-11</td><td>All data returned by the NVMe Basic Management Command shall be returned in big-endian format unless otherwise noted.</td></tr>
</table>

11.2 NVMe Basic Management Command (Appendix A) Data Format

<table>
<tr><th>Command Code (Decimal)</th><th>Byte Offset (Decimal)</th><th>Description</th></tr>
<tr><td>0</td><td>0</td><td>Defined in the NVM Express Management Interface Specification.</td></tr>
<tr><td>8</td><td>8</td><td>Defined in the NVM Express Management Interface Specification.</td></tr>
<tr><td>32</td><td>32</td><td>Payload length of Command Code 32: This is the number of bytes until the PEC code. This shall be set to 10h.</td></tr>
<tr><td></td><td>48:33</td><td>GUID: This is a 16-byte Global Unique Identifier:
<table>
<tr><th>Byte Address</th><th>Value</th></tr>
<tr><td>33</td><td>73h</td></tr>
<tr><td>34</td><td>89h</td></tr>
<tr><td>35</td><td>20h</td></tr>
<tr><td>36</td><td>E5h</td></tr>
<tr><td>37</td><td>6Bh</td></tr>
<tr><td>38</td><td>EEh</td></tr>
<tr><td>39</td><td>42h</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Command Code (Decimal)</th><th>Byte Offset (Decimal)</th><th>Description</th></tr>
<tr><td></td><td></td><td><table>
<tr><td>40</td><td>58h</td></tr>
<tr><td>41</td><td>9Ah</td></tr>
<tr><td>42</td><td>7Ah</td></tr>
<tr><td>43</td><td>CEh</td></tr>
<tr><td>44</td><td>BDh</td></tr>
<tr><td>45</td><td>B3h</td></tr>
<tr><td>46</td><td>5Fh</td></tr>
<tr><td>47</td><td>00h</td></tr>
<tr><td>48</td><td>85h</td></tr>
</table></td></tr>
<tr><td></td><td>49</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td>50</td><td>50</td><td>Payload length of Command Code 50: Indicated the number of additional bytes to read before encountering PEC. Shall be set to 26h.</td></tr>
<tr><td></td><td>51</td><td>Temperature Flags: This field reports the effect of temperature on the device’s performance.
Temperature Throttling – Bit 7 is set to 1b when the device is throttling performance to prevent overheating. Clear to 0b when the device is not throttling.
Bits 6:0 shall be set to 1111111b.</td></tr>
<tr><td></td><td>52</td><td>Max Power Supported: This shall denote the Max Average Power (MAP) supported by this device rounded to the nearest watt (e.g., a 50W device is 32h, a 25W device is 19h, a 15W device is 0Fh, an 8.25W device is 8W which is 08h).</td></tr>
<tr><td></td><td>84:53</td><td>Configured NVMe Power State: This is a copy of the NVMe Power State Descriptor Data Structure of the currently configured Power State and is laid out in little-endian format.</td></tr>
<tr><td></td><td>88:85</td><td>Total NVM Capacity: This field indicates the total usable NVM capacity in the NVM subsystem in GB in Hex (2048 GB in total capacity = 0000800h). This field is equivalent to the TNVMCAP field in the Identify Controller Data Structure.</td></tr>
<tr><td></td><td>89</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td>90</td><td>90</td><td>Payload length of Command Code 90: Indicates number of additional bytes to read before encountering PEC. Shall be set to 04h.</td></tr>
</table>

<table>
<tr><th>Command Code (Decimal)</th><th>Byte Offset (Decimal)</th><th>Description</th></tr>
<tr><td rowspan="3"></td><td>91</td><td>Firmware Update Flags: This field allows the host to control whether the current firmware allows new firmware images to be activated (see Section 12 Security for more information).
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7</td><td>Written by host, read by device:
<table>
<tr><td>1b</td><td>Unlock Firmware Update
Device shall enable Firmware update</td></tr>
<tr><td>0b</td><td>Lock Firmware Update
Device shall block and error on Firmware download and activate commands</td></tr>
</table></td></tr>
<tr><td>6</td><td>Written by device, read by host:
<table>
<tr><td>1b</td><td>Firmware Update Unlocked
Device shall allow Firmware download and activate commands</td></tr>
<tr><td>0b</td><td>Firmware Update Locked
Device shall block and error on Firmware download and activate commands</td></tr>
</table></td></tr>
<tr><td>5:0</td><td>Shall be set to 111111b.</td></tr>
</table>
The device shall revert to the default Unlock/Lock state on the next power cycle. If the host attempts to issue a Download Firmware command when the device is in the Firmware Update Locked state, the device shall fail the command with status Operation Denied.</td></tr>
<tr><td>94:92</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>95</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td rowspan="3">96</td><td>96</td><td>Payload length of Command Code 96: Indicates number of additional bytes to read before encountering PEC. Shall be set to 38h.</td></tr>
<tr><td>104:97</td><td>Firmware Version Number: This field shall indicate the activated firmware version that is running on the device after the firmware activation took place. The format of this field shall be as defined in field Firmware Revision (FR) Section 5.15.2.2 Identify Controller Data Structure of the NVMe specification version 1.4b.</td></tr>
<tr><td>112:105</td><td>Security Version Number: This is the Security Version Number of the currently running firmware image. The supplier increments this number any time a firmware includes a fix for a security issue or critical firmware fix that customer agrees rollback prevention is required. This is a copy of SMART-22 – Security Version Number.</td></tr>
</table>

<table>
<tr><th>Command Code (Decimal)</th><th>Byte Offset (Decimal)</th><th>Description</th></tr>
<tr><td></td><td>152:113</td><td>Model Number: This shall be a copy of the data in the Model Number field of Identify Controller Data Structure (CNS 01h, byte offset 63:24).</td></tr>
<tr><td></td><td>153</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in SMBus Specification.</td></tr>
<tr><td>154</td><td>154</td><td>Payload length of Command Code 154: Indicates number of additional bytes to read before encountering PEC. Shall be set to 0Bh.</td></tr>
<tr><td></td><td>155</td><td>Panic Rest Action: See EREC-2.</td></tr>
<tr><td></td><td>156</td><td>Device Recovery Action 1: See EREC-3.</td></tr>
<tr><td></td><td>164:157</td><td>Panic ID: See EREC-4.</td></tr>
<tr><td></td><td>165</td><td>Device Recovery Action 2: See EREC-14.</td></tr>
<tr><td></td><td>166</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td>167</td><td>167</td><td>Payload length of Command Code 167: Indicates number of additional bytes to read before encountering PEC. Shall be set to 20h.</td></tr>
<tr><td></td><td>199:168</td><td>Panic Context Buffer: Vendor may record additional information about the Panic ID reported in Command Code 154. Any unused bytes shall be cleared to zero.</td></tr>
<tr><td></td><td>200</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td>201</td><td>241:201</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td>242</td><td>242</td><td>Payload length of Command Code 242: Indicates number of additional bytes to read before encountering PEC. Shall be set to 04h.</td></tr>
<tr><td></td><td>243</td><td>Secure Boot Failure Reporting:
<table>
<tr><th>Bit</th><th>Bit Description</th></tr>
<tr><td>7</td><td>Secure Boot Failure Feature Reporting Supported:
When set to 1b the secure boot feature reporting is supported. When cleared to 0b the secure boot failure feature reporting is not supported.</td></tr>
<tr><td>6</td><td>Secure Boot Failure Status: When cleared to 0b there is no secure boot failure. When set to 1b there is a secure boot failure. This bit shall only be set if the Secure Boot Feature Supported bit is set to 1b and there is a secure boot failure.</td></tr>
<tr><td>5</td><td>OCP Recovery/ Platform Root-of-Trust for Recovery: When set to 1b, OCP Recovery/ Platform Root-of-Trust for Recovery codes are supported in byte 244. When cleared to 0b OCP Recovery/</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Command Code (Decimal)</th><th>Byte Offset (Decimal)</th><th>Description</th></tr>
<tr><td></td><td></td><td><table><tr><td></td><td>Platform Root-of-Trust for Recovery codes are not supported and byte 244 shall be cleared to zero.</td></tr><tr><td>4:0</td><td>Reserved. Shall be cleared to zero.</td></tr></table></td></tr>
<tr><td></td><td>244</td><td>Recovery Code: OCP Recovery/Platform Root-of-Trust for Recovery code.</td></tr>
<tr><td></td><td>246:245</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td></td><td>247</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
<tr><td>248</td><td>248</td><td>Payload length of Command Code 248: Indicates number of additional bytes to read before encountering PEC. Shall be set to 06h.</td></tr>
<tr><td></td><td>250:249</td><td>Data Format Version Number: Indicates the version of this mapping used in the device. Shall be set to 0004h.</td></tr>
<tr><td></td><td>254:251</td><td>Reserved. Shall be cleared to zero.</td></tr>
<tr><td></td><td>255</td><td>PEC: An 8-bit CRC calculated over the SMBus address, command code, second SMBus address and returned data. The algorithm is defined in the SMBus Specification.</td></tr>
</table>

11.3 NVMe-MI Requirements

Requirement ID	Description
NVMe-MI-1	The device shall support NVMe Management Interface Specification version 1.2c or later, including all mandatory requirements. The vendor shall provide an NVMe-MI compliance report that demonstrates compliance.
NVMe-MI-2	The device shall support all mandatory and optional NVMe-MI functionality over SMBus (see NVMe-MI-1), even when the PCIe link is not active (as required by the NVMe Management Interface Specification).
NVMe-MI-3	The device shall support NVMe-MI over PCIe VDM (see NVMe-MI-1).
NVMe-MI-4	The device shall support SMBus Fixed and Discoverable.
NVMe-MI-5	The device shall NACK any SMBus addresses not listed in the NVM Express Management Interface Specification.
NVMe-MI-6	The device shall support the GET UDID command.
NVMe-MI-7	The device shall have a VPD that is accessible over SMBus and supports IPMI Platform Management FRU Information. The VPD shall support all required elements that are defined in the NVM Express Management Interface Specification.
NVMe-MI-8	Obsolete, see NVMe-MI-2.
NVMe-MI-9	Obsolete, see NVMe-MI-3.
NVMe-MI-10	All mandatory NVMe Admin commands and the following optional NVMe Admin commands shall be supported by MCTP over SMBus, and by MCTP over PCIe VDM.

Requirement ID	Description
	• Firmware Activate/Commit • Firmware Image Download • Lockdown • Sanitize • Security Send for TCG devices • Security Receive for TCG devices • Device Self-test • Set Features
NVMe-MI-11	All mandatory Log Identifiers and the following optional Log Identifiers shall be supported by MCTP over SMBus, and by MCTP over PCIe VDM. • Device Self-test (06h) • Command and Feature Lockdown (14h) • Sanitize Status (81h) • SMART / Health Information Extended (C0h)
NVMe-MI-12	The following Log Identifiers shall be supported by MCTP over SMBus, and by MCTP over PCIe VDM. • Persistent Event Log (0Dh) • Telemetry Host-Initiated (07h) • Telemetry Controller-Initiated (08h)
NVMe-MI-13	The following Feature Identifiers shall be supported by MCTP over SMBus, and by MCTP over PCIe VDM. • Power Management (02h) • Timestamp (0Eh) • Temperature Threshold (04h)
NVMe-MI-14	The device shall support locking down all supported Admin commands and all supported Features on the in-band interface via a Lockdown command issued out-of-band via MCTP over SMBus, and by MCTP over PCIe VDM. This includes the Lockdown command itself and all Admin commands and Features not referenced in this specification.
NVMe-MI-15	The following commands are prohibited and shall return an NVMe-MI Response Message Status of Invalid Command Opcode: • PCIe Configuration Write • PCIe Memory Write • PCIe Memory Read • PCIe I/O Write • PCIe I/O Read Support for the PCIe Configuration Read command is optional.
NVMe-MI-16	Access to user data is prohibited via out-of-band NVMe-MI (e.g., vendor-specific commands shall not access user data).
NVMe-MI-17	The device shall support SMBus Frequencies of 100 KHz and 400 KHz.

Requirement ID	Description
NVMe-MI-18	For SMBus, the device shall support all possible MCTP Transmission Unit Size values up to the largest (250 bytes).
NVMe-MI-19	For PCIe VDM, the NVMe device shall support all possible MCTP Transmission Unit Size values up to the PCIe Max Payload Size Supported.
NVMe-MI-20	NVMe devices shall support a minimum throughput of 50 MiB/s for reads and writes for the MCTP over PCIe VDM interface under the assumption that there is no delay between Request Messages from the Management Controller for transfers that require multiple Request/Response Messages.
NVMe-MI-21	Any TCG method invoked via NVMe-MI that takes longer than 100ms shall run in the background using the OutstandingData mechanism.
NVMe-MI-22	The device shall support NVMe-MI over SPDM Secured Messages (see SPDM-7).

12 Security

12.1 Basic Security Requirements

Requirement ID	Description
SEC-1	The device shall support signed firmware binary update which is checked before firmware is activated. The device firmware shall be authenticated using cryptographic keys on every reboot and during firmware update.
SEC-2	The device shall support XTS-AES-256 (as specified in NIST SP 800-38E, which references IEEE STD 1619) or AES-256-GCM (as specified in NIST SP 800-38D) hardware-based data encryption for performance and to prevent firmware tampering with the encryption mechanism. AES-256-GCM is the preferred mode. The combination of AES-256-CTR (as specified in NIST SP 800-38A) with HMAC-SHA256 (where the HMAC is calculated on the ciphertext) is an acceptable alternative.
SEC-3	The device shall support anti-rollback protection for firmware. The anti-rollback protection shall be implemented with a security version which is different than the firmware version. If the security version of the firmware being activated is greater or equal to the current security version the firmware may be activated. If the security version of the firmware being activated is not equal or greater than the firmware being activated the firmware update shall fail.
SEC-4	The device shall support Crypto Erase (see NVMe-AD-5) and NVMe-AD-7).
SEC-5	The device shall support Secure Boot (see Section 12.2 Secure Boot).
SEC-6	The device shall have a method of identifying a secure boot failure which does not require physical access to the device.
SEC-7	The device’s cryptographic module shall be FIPS 140-3 level 2 capable per CMVP and shall follow the NIST 800-90 (A, B and C) specification.
SEC-8	The device shall only implement NIST approved cryptographic algorithms and shall be designed to achieve FIPS-140-3 validation (see SEC-7).

Requirement ID	Description
SEC-9	The device shall support key revocation allowing a new key to be used for firmware verification on update. The preferred implementation is to allow for up to 3 key revocations.
SEC-10	Obsolete. Moved to TCG-FEATURE-1.
SEC-11	Obsolete. Moved to TCG-FEATURE-2.
SEC-12	Obsolete. Moved to TCG-FEATURE-3.
SEC-13	For models that support eDrive, the IEEE 1667 Probe (0x100) and TCG Storage Transport (0x104) silos shall be supported.
SEC-14	Supplier shall follow the Security Development Lifecycle (SDL), and provide a report with the following for each qualification-ready or production-ready firmware version: • The Threat Model. • Fuzz & Pen Tests. • Static Analysis. • Build Logs and Compiler Security Settings. Additional information about the SDL is available here: *https://www.microsoft.com/en-us/sdl/default.aspx*
SEC-15	Security audits, including firmware source code review, shall be provided to the customer. This will include Telemetry and debug logs, etc.
SEC-16	All signing keys shall be stored in a Hardware Security Module (HSM) that is either FIPS 140-2 Level 3 (or greater) certified or FIPS 140-3 Level 3 (or greater) certified.
SEC-17	Access/use of signing keys should be restricted to a small set of developers, following the principle of least privilege. The number of people with access and their corresponding roles shall be provided.
SEC-18	All debug ports shall be disabled before the device leaves the factory. Alternatively, the ports shall only be accessible in the field after a successful exchange of a challenge-response mechanism using an asymmetric crypto scheme (see NIST SP 800-63). The state shall be reset to inaccessible on any reset or power cycle.
SEC-19	All vendor unique commands, log pages or set features that are not explicitly defined in this specification or approved of in writing by the customer shall be disabled before the device leaves the factory. Alternatively, the commands/log pages/set features shall only be accessible in the field after a successful exchange of a challenge-response mechanism using an asymmetric crypto scheme (see NIST SP 800-63). The state shall be reset to inaccessible on any reset or power cycle.
SEC-20	Adversarial testing using red teams shall be conducted before the start of customer qualification of the initial SSD product and any major firmware or hardware revision. A report of items attempted, and results shall be provided to the customer.
SEC-21	The vendor shall provide timely notification to the customer of security issues and delivery of fixes to the customer: • The vendor shall document all security fixes with each firmware update. • The vendor shall notify end customer within 7 days of discovering security issues in the device hardware or firmware.

Requirement ID	Description
	● Notification of issues shall include the process and timeline of the vendor’s commitment to fix the issue: ○ For privately disclosed vulnerabilities, the duration shall be no longer than 90 days. ○ For publicly disclosed vulnerabilities, the duration shall be no longer than 7 days. ○ Vendors shall notify the customers about the known CVEs and security issues and provide security-related updates before public announcement.
SEC-22	All Telemetry and debugging logs shall be designed to be human readable (e.g., made available to the customer in a self-describing JSON format via a vendor-supplied tool), except for Telemetry Data Areas 3 and 4.
SEC-23	The device shall not include user data, passwords, keys or any secret or sensitive information in any Telemetry or debug logs.
SEC-24	All public keys shall be revocable.
SEC-25	Secure Boot Flow shall be based upon a hardware root of trust. All mutable Key(s), Certificate(s) and/or firmware shall be cryptographically bound to the hardware root of trust. See Section 12.2 Secure Boot.
SEC-26	The device shall only update firmware via a secure firmware update flow that shall be based upon immutable public keys.
SEC-27	Obsolete. Moved to TCG-CRYPTO-1.
SEC-28	The device shall delete all keys from volatile memory as soon as they are no longer needed for operation during the current power on state.
SEC-29	The device shall only store host provided passwords, host provided keys, or any host provided secret information in non-volatile memory at any stage in an encrypted form. The encryption key for this protection shall not be stored in non-volatile memory. Device shall not store plaintext/cleartext secrets in any non-volatile memories.
SEC-30	The supplier must deliver key management and encryption flow diagram details (not source code) which includes: ● Encryption algorithms and modes (e.g., RNG, wrapping function, key derivation function). ● Key size and password length. ● Any crypto information that is stored in nonvolatile memory. ● Critical Security Parameters. Complete details about algorithm inputs (e.g., Initialization vector source and size, salt source and size, unique per device/product line/vendor)
SEC-31	The supplier must provide industry certification reports, if available, such as FIPS, NIST for device and device components such as TRNG, RNG, Crypto engine etc.
SEC-32	Log data and user data (data transferred from the Host in Write Commands) shall be stored on physically separate areas on the device. For example, in the system area and the user data area, respectively.
SEC-33	The vendor shall provide a comprehensive list of what is included in the logs covering both the standard logs and vendor-unique logs. The list shall identify any defined optional events or items that are not logged by the device.

Requirement ID	Description
SEC-34	Obsolete.
SEC-35	All telemetry and debug logs shall only be writable by device firmware that has been verified as part of a secure boot.
SEC-36	Obsolete. Moved to TCG-AUTH-3.
SEC-37	Obsolete. Moved to NVME-MI-2 and NVME-MI-10.
SEC-38	Obsolete. Moved to TCG-FEATURE-4.
SEC-39	The device shall not allow the host to read firmware or boot code from the device.
SEC-40	All device firmware packages shall be encrypted based on a symmetric encryption algorithm. See SEC-2 for acceptable algorithms and SEC-29 for key storage requirements.
SEC-41	The device shall only support activation of encrypted firmware images.
SEC-42	The device shall only store firmware in encrypted format when stored in non-volatile memory.
SEC-43	Sanitize, Format NVM User Data Erase/Cryptographic Erase, and TCG Revert/RevertSP/GenKey shall comply with the IEEE 2883-2022 Purge requirements.
SEC-44	Any command that changes the media encryption key(s) as part of its operation (e.g., Sanitize command with Crypto Erase (100b), TCG Revert, TCG RevertSP, etc.), shall purge, as defined in IEEE 2883-2022, all copies of the previous key(s) before the device completes the operation. If the device cannot guarantee that all copies of the previous media encryption key(s) have been purged, then the device shall fail the operation.
SEC-45	Any optional security features supported by the device and not described in this document shall be clearly documented and disclosed to the customer.
SEC-46	TCG devices shall support the Security Receive and Security Send commands in-band (i.e., bit 0 in the Optional Admin Command Support field in the Identify Controller data structure shall be set to 1b).

12.2 Secure Boot

The device shall support Secure Boot. There are two fundamental things to address for secure boot:

- Secure boot rooted in hardware.
- Core Root of Trust Measurement.

The vendor should follow the recommendations in the *TCG Hardware Requirements for a Device Identifier Composition Engine* and the guidelines in NIST SP 800-193. DICE coupled with *TCG Implicit Identity Based Device Attestation* and Source for *RIoT Reference Architecture* can help implement Cryptographic Identity with implicit attestation.

Requirement ID	Description
SBT-1	The device shall comply with the *FIPS 186-5 Digital Signature Standard (DSS)* and the *OCP Hardware Secure Boot V1.0*. This requirement does not require the full implementation of the OCP Recovery specification, only the Recover Error Codes (See SMBUS-6). This requirement does not require implementation of the OCP Attestation of System Components specification.

Requirement ID	Description
SBT-2	For Core Root of Trust measurement, each device shall have a Cryptographic Device Identity.
SBT-3	The *TCG Hardware Requirements for a Device Identifier Composition Engine* standard, or hardware based cryptographic identity shall be implemented.
SBT-4	The device shall follow the guidance in the *Commercial National Security Algorithm Suite 2.0* regarding quantum resistant algorithms, key sizes, and transition timelines.
SBT-5	Secure boot flow shall be immutable for exploitation and use immutable public keys.

12.3 DMTF Security Protocol and Data Model (SPDM)

The vendor should follow the guidelines in NIST SP 800-193.

Requirement ID	Description
SPDM-1	The device shall support firmware measurement and identity authentication per the DMTF SPDM 1.2 or later specifications. The device shall be backwards compatible with SPDM 1.1 and later specifications.
SPDM-2	The device shall support the PCI-SIG Component Measurement and Authentication (CMA) ECN.
SPDM-3	The device shall support SPDM-1 and SPDM-2 via MCTP over SMBus/I2C.
SPDM-4	The device shall support SPDM-1 and SPDM-2 via MCTP over PCIe VDM.
SPDM-5	The device shall support SPDM-1 and SPDM-2 via PCI-SIG Data Object Exchange (DOE).
SPDM-6	The leaf certificate shall include a Subject Alternative Name extension as defined in the CMA ECN.
SPDM-7	The device shall support Secured Messages using SPDM over MCTP Binding Specification per the DMTF DSP0276 Version 1.0 or later.

12.4 Data Encryption and Eradication

Requirement ID	Description
DATAE-1	Obsolete. Duplicate of SEC-2.

12.5 TCG Implementation Requirements

Devices that support TCG functionality (i.e., TCG devices) shall implement the specifications in accordance with the additional requirements in this section.

Feature/Parameters		Value
SSC Feature Descriptor		Opal 2 (0x0203)
Core Spec (TCG Storage Architecture)		2.01
SID	# of SIDs	1
	TryLimit	10
	Persistence	FALSE
	Admin authorities	4 (min)

Feature/Parameters		Value
Number of Locking SP	User authorities	8 (min)
Number of Locking Ranges	Global	1
	Additional	8 (min)
Data Store size		10 MB (min)
MBR Table		Optional*
Configurable Access Control		Mandatory
Life Cycle support		Mandatory
Repurpose (Revert/ Revert SP)		Mandatory
Shadow MBR		Prohibited*
Feature Set	Single User Mode	Mandatory
	PSK Secure Messaging	Optional
	PSID	Mandatory
	Block SID	Mandatory
	Additional Data Store Tables	Mandatory 5 (min) @ 128KB each (not included in Data Store size)

* This overrides TCG specification requirements for this item.

12.5.1 TCG Version and Features

Requirement ID	Description
TCG-FEATURE-1	TCG devices shall support TCG Opal v2.02 or later, and all mandatory feature sets (e.g., Additional Datastore and PSID feature sets).
TCG-FEATURE-2	TCG devices shall support TCG Single User Mode feature set Version 1.00, revision 2.00 or later.
TCG-FEATURE-3	TCG devices shall support TCG Configurable Namespace Locking (CNL) feature set Version 1.00, revision 1.00 with mandatory support for the Namespace Global Range Locking object. The Namespace Non-Global Range Locking object may be supported.
TCG-FEATURE-4	TCG devices shall support TCG Feature Set: Block SID Authentication.
TCG-FEATURE-5	The current TCG feature set configuration settings, if any, shall persist through a download of firmware code. This is required regardless of the "default setting state" of the new code being downloaded.

12.5.2 TCG Communication Layer

Requirement ID	Description
TCG-COMM-1	The device shall support ComID management, and the Level 0 Discovery TPer Feature Descriptor shall report ComID Mgmt supported.

Requirement ID	Description
TCG-COMM-2	The device shall support at least 2 Static ComIDs.
TCG-COMM-3	The device should support at least 2 Dynamic ComIDs.
TCG-COMM-4	MaxComIDTime shall be at least 120 seconds.
TCG-COMM-5	MaxComIDTime shall be greater than DefSessionTimeout.

12.5.3 TCG Session Layer

Requirement ID	Description
TCG-SES-1	A single Read-Write Session shall be supported per SP and shall be indicated by MaxSessions. (It is not required to support an active RW session on the AdminSP at the same time as sessions on other SPs, because this would violate the TCG Core spec.)
TCG-SES -2	At least 2 Read-Only Sessions shall be supported per SP and shall be indicated by MaxReadSessions.
TCG-SES -3	DefSessionTimeout shall be capable of being read and written using the Properties Method.
TCG-SES -4	The default DefSessionTimeout shall be 90,000 milliseconds.
TCG-SES -5	MaxSessionTimeout shall be 0. This value indicates that there is no limit.
TCG-SES -6	MinSessionTimeout should have a value of 500 (units in milliseconds).
TCG-SES -7	SessionTimeout shall be supported in the StartSession method call.
TCG-SES -8	Support for setting the SPSessionTimeout column of the SPInfo table is recommended.

12.5.4 TCG Transactions

Requirement ID	Description
TCG-TRAN-1	TCG devices shall support TCG Transactions.
TCG-TRAN-2	TCG Nested Transactions may be supported. If TCG Nested Transactions are supported, all supported transactions including any nested transactions shall be reflected in the MaxTransactionLimit property.

12.5.5 TCG Authentication

Requirement ID	Description
TCG-AUTH-1	TCG devices shall enforce a minimum password length (C-PIN PinLength of 16) of 16 bytes for all TCG Opal authorities.
TCG-AUTH-2	TCG devices shall allow a password length (C-PIN PinLength of 32) of at least 32 bytes for all TCG Opal authorities.
TCG-AUTH-3	TCG devices shall set the default TryLimit to 10 for all TCG Opal authorities. The value of the Tries column shall only be cleared to 0 for the mandatory reasons specified in the TCG Storage Architecture Core Specification (e.g., do not clear the Tries column via the Set method, do clear it upon successful authentication or power cycle).

Requirement ID	Description
TCG-AUTH-4	Authentication of a C_PIN shall be delayed limiting how fast the next authentication can occur. • The delay shall be at least 100ms and should not exceed 250ms. • The delay shall apply regardless of whether the authentication is successful. • The device manufacturer shall supply documentation with the implemented delay value. • The time limit should be enforced cryptographically (e.g., by tuning the number of PBKDF iterations to reach the target time).

12.5.6 TCG Crypto

Requirement ID	Description
TCG-CRYPTO-1	If TCG Opal SP has been activated and the host has supplied a locking range password, the device shall use that password as an added source of randomness for its Deterministic Random Number Generator (DRNG). For example, the device may use KDF (PBKDF2) on the user provided locking range password to generate an output of equal length to the required locking range key length plus DRNG seed length. This output can then be utilized in parts, one part as Opal locking range key and the other part as an additional input to the DRNG. If the DRNG implementation does not allow the use of additional input, then an alternate option would be to XOR this output (e.g., from PBKDF2 operation) with initial DRNG seed.

12.5.7 TCG Locking Security Provider

Requirement ID	Description
TCG-LOCKING-1	The number of Locking Ranges supported shall be greater than or equal to the maximum number of NVMe namespaces supported on the device.
TCG-LOCKING-2	The device shall support at least one User per Locking Range.
TCG-LOCKING-3	The K_AES_256 Table shall be supported.
TCG-LOCKING-4	The K_AES_256 Table’s Key columns shall not be readable by the host or be accessible outside the device in any manner.
TCG-LOCKING-5	Any attempt to remove the Power Cycle reset type from a LockOnReset column set value shall be rejected.

12.5.8 TCG PSID Feature Set

Requirement ID	Description
TCG-PSID-1	The PSID shall not be deterministic, shall comprise at least 160 bits of entropy. The entropy shall be unique for each device. The PSID shall not be derived from any information that is available without physical access to the device.

Requirement ID	Description
	For example, the PSID shall not be the same as the MSID, the PSID shall not be derived from the device serial number or the date of manufacture because these can be read via the device interface or obtained from other documentation. The PSID shall not be the serial number of the drive hashed with a secret string shared with any other device or among all devices, because the entropy shall be unique for each device. The presence of the PSID on the device label does not violate this requirement.
TCG-PSID-2	The PSID shall consist of 32 ASCII characters. The allowed ASCII characters are upper case letters, lower case letters and the digits 0-9. The characters 'I', 'l', 'O' and 'o' (i.e., 49h, 6Ch, 4Fh and 6Fh) and the digits '0' and '1' (i.e., 30h and 31h) shall not be used to limit possible visual confusion.
TCG-PSID-3	The generation method for the PSID shall be documented and provided to the customer.
TCG-PSID-4	After the PSID is provisioned in manufacturing, the PSID shall not be persistently stored in a form that makes retrieval or reconstruction of the PSID by firmware possible, either on the drive or outside the drive (e.g., it is acceptable to store the PSID on the drive hashed via a cryptographic hash or KDF such that the hashed PSID can only be used to verify a PSID input from the host). The presence of the PSID on the device label does not violate this requirement.
TCG-PSID-5	Once manufacturing is complete, the PSID shall not be stored anywhere other than on the physical drive label.

13 Device Profiles

The following are device profiles. This section is intended to be firmware-based configuration settings configured by device suppliers when manufacturing a device. A device may be configured with a mix of A and/or B settings. Each customer shall provide their A/B preference for each configuration setting.

The following conventions are used for the Device Profile Table:

Convention	Definition
R	Required. This shall be supported.
O	Optional. This may be supported.
P	Prohibited. This shall not be supported.

Requirement ID	Description	Configuration Setting	
		A	B
DP-CFG-1	Factory Default Sector Size.	4096-byte	512-byte
DP-CFG-2	Number of Namespaces Supported.	NSM-4 (16 Namespaces)	NSM-5 (16 Namespaces per TB)
DP-CFG-3	Retention Time based on RETC-1.	1 Month	3 Months
DP-CFG-4	NVMe Basic Management Command Supported.	R	P

Requirement ID	Description	Configuration Setting	
		A	B
DP-CFG-5	Max M.2 top side height.	2.0mm	3.2mm
DP-CFG-6	EOL/PLP Failure Mode (Feature Identifier C2h).	Enabled	Disabled
DP-CFG-7	Write Uncorrectable command support.	O	R
DP-CFG-8	Time-to-Identify-Ready based on TTR-1.	<= 1 second	<= 10 seconds after a normal shutdown <= 60 seconds after an abrupt shutdown
DP-CFG-9	Time-to-I/O-Ready based on TTR-2.	<= 20 seconds	<= 10 seconds after a normal shutdown <= 60 seconds after an abrupt shutdown
DP-CFG-10	In addition to the requirements in TTR-4, the device shall keep CSTS.RDY = 0 until the device is able to service I/O commands successfully.	P	R
DP-CFG-11	Support for the Lockdown command on the in-band interface.	R	P

14 Labeling

The following sample label is meant to be used to refer to the label requirements in Section 14.1 – Label Requirements. It is not a model label and any markings on it are informative only. See the specifics in Section 14.1 – Label Requirements for actual requirements:

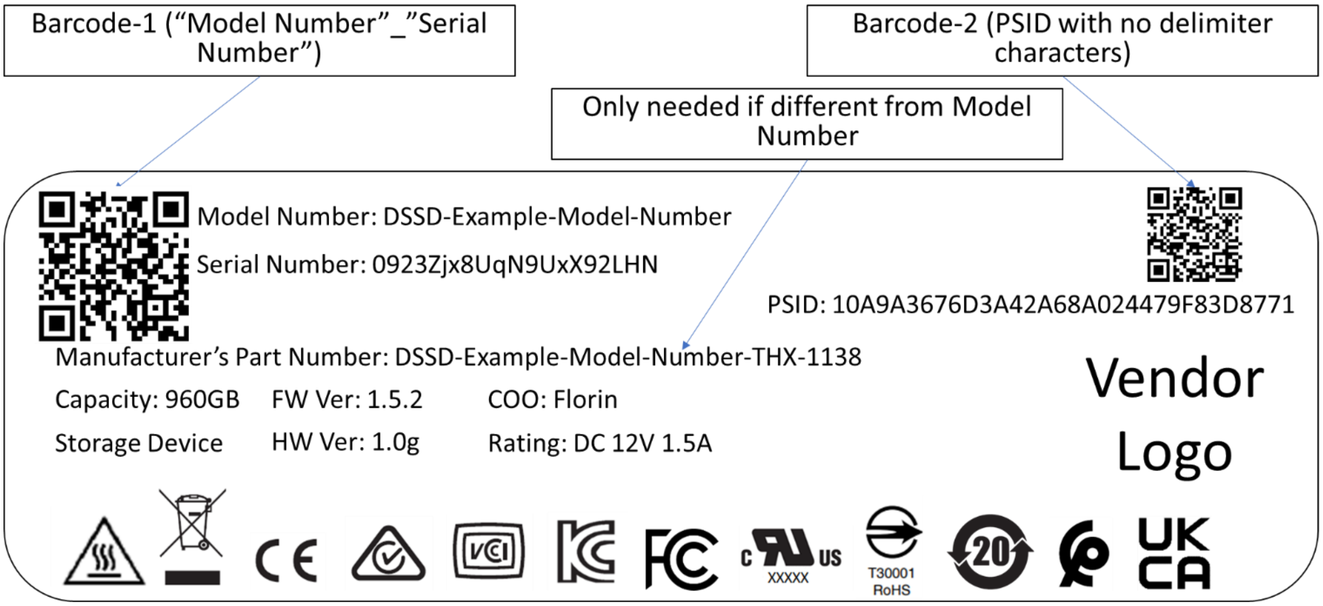
Barcode-1 (“Model Number”_”Serial Number”)
Barcode-2 (PSID with no delimiter characters)
Only needed if different from Model Number
Model Number: DSSD-Example-Model-Number
Serial Number: 0923Zjx8UqN9UxX92LHN
PSID: 10A9A3676D3A42A68A024479F83D8771
Manufacturer’s Part Number: DSSD-Example-Model-Number-THX-1138
Capacity: 960GB
FW Ver: 1.5.2
COO: Florin
Storage Device
HW Ver: 1.0g
Rating: DC 12V 1.5A
Vendor Logo
XXXXX
T30001
RoHS
20
UK CA

14.1 Label Requirements

Requirement ID	Description
LABL-1	The following fields are required information that shall be placed on the label:

Item	Format	Text Required	Barcode Required	Barcode Type
Barcode-1	‘Model Number’ ‘Underscore’ ‘Serial Number’ \n.	No	Yes	2d
Model Number	See LABL-11 (Model Number shall match).	Yes	No	N/A
Serial Number	See LABL-12 (Serial Number shall match), LABL-15 (certification logos), and LABL-17 (Serial Number format).	Yes	No	N/A
Manufacturer’s Part Number	Number used for ordering.	Yes, if different from Model Number	No	N/A
Capacity	Number of GB or TB.	Yes	No	N/A
STORAGE DEVICE	Text shall be “STORAGE DEVICE”.	Yes	No	N/A
PSID	TCG-OPAL Spec.	Yes	No	N/A
Barcode-2	‘PSID’ \n	No	Yes	2d
HW Revision		Yes	No	N/A
Firmware Name		No	No	N/A

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><td>& Revision</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Regulatory Mark</td><td colspan="2">See LABL-15 (certification logos).</td><td>No</td><td>N/A</td></tr>
<tr><td>Country Certification Numbers</td><td colspan="2">See LABL-15 (certification logos).</td><td>No</td><td>N/A</td></tr>
<tr><td>Certification Logos</td><td colspan="2">See LABL-15 (certification logos).</td><td>No</td><td>N/A</td></tr>
<tr><td>RoHS/ Green</td><td colspan="2">See LABL-15 (certification logos).</td><td>No</td><td>N/A</td></tr>
</table>
</td></tr>
<tr><td>LABL-2</td><td>The Model Number on the shipping label shall match the Model Number used during qualification.</td></tr>
<tr><td>LABL-3</td><td>The minimum font size shall be 3 points and the typical size should be 6 points.</td></tr>
<tr><td>LABL-4</td><td>For the Capacity field, if there are space constraints, the manufacturer may remove “Capacity:” and just show “960GB” or “60TB” for example.</td></tr>
<tr><td>LABL-5</td><td>To distinguish Model Number and Serial Number, Barcode-1 shall have an underscore “_” between the Model Number portion and the Serial Number portion.
Example:
Model Number: DSSD-Example-Model-Number
Serial Number: 0923Zjx8UqN9UxX92LHN
Barcod-1 Readout: DSSD-Example-Model-Number_0923Zjx8UqN9UxX92LHN</td></tr>
<tr><td>LABL-6</td><td>There shall be a line with the text “STORAGE DEVICE”.</td></tr>
<tr><td>LABL-7</td><td>The following fields are optional information that can be placed on the label at the discretion of the device maker. Placement is also at the device makers’ discretion if such information does not interfere with the mandatory information above. No additional barcode shall be present.
<table>
<tr><th>Item</th><th>Format</th><th>Text Required</th><th>Barcode Required</th><th>Barcode Type</th></tr>
<tr><td>Processor Code (BA)</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Maker Logo</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Rated Voltage & Current</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Production Date</td><td>DDMMYYYY: DD (Date), MM (Month), YYYY (Year)</td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Weekly Code</td><td>YYWW: YY (Year), WW (Week)</td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Warranty VOID IF REMOVED</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><td>Makers Own Label Material Number</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Website, Company Address</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>SSD</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>Product Series Name</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>SA: Value used within manufacturing</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>PBA: Physical Board Address (identifies the physical configuration of the device)</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
<tr><td>WWN: World Wide Number (unique for each device)</td><td></td><td>Optional</td><td>No</td><td>N/A</td></tr>
</table>
</td></tr>
<tr><td>LABL-8</td><td>To ensure that datacenter operations personnel can quickly and easily identify devices that have been ticketed for field replacement, it is mandatory to have the proper identifying fields on the label(s), in the format specified below.</td></tr>
<tr><td>LABL-9</td><td>The label shall not degrade over the standard SSD lifetime under standard operating conditions.</td></tr>
<tr><td>LABL-10</td><td>For each form factor, the label shall be placed as specified below:
• M.2: the label shall be placed on the top side of the device as defined in the PCI-SIG M.2 form factor specification.
• E1.S: the label shall be placed on the Primary side of the device as defined in the SFF TA-1006 form factor specification.
• E1.L: the label shall be placed on either the Primary or Secondary side of the device as defined in the SFF TA-1007 form factor specification.
• E3: the primary label shall contain the drive information (e.g., model number, serial number, manufacturing info, firmware revision) and be placed on the top label area as defined in the SFF-TA-1008 form factor specification. If more space is needed a label may also be placed on the bottom label area with regulatory marks.</td></tr>
<tr><td>LABL-11</td><td>The Model Number in Barcode-1, the Model Number printed on the label and the Model Number returned in Identify Controller Data Structure (CNS 01h, byte offset 63:24) shall all match unless exempted in writing by the customer.</td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>LABL-12</td><td>The Serial Number in Barcode-1, the Serial Number printed on the label and the Serial Number returned in Identify Controller Data Structure (CNS 01h, byte offset 23:04) shall all match.</td></tr>
<tr><td>LABL-13</td><td>This Hardware Revision printed on the label and returned by the NVMeCLI utility shall match.</td></tr>
<tr><td>LABL-14</td><td>All other electronically readable information shall also match their counterparts printed on the label.</td></tr>
<tr><td>LABL-15</td><td>The following certification logos and their corresponding certifications are required:
<table>
<tr><th>Regulatory Mark/Text</th><th>Description</th></tr>
<tr><td>Regulatory Model Number</td><td>Unique regulatory Identifier.</td></tr>
<tr><td>Made in</td><td>Country of Origin.</td></tr>
<tr><td>Manufacturer or Brand name</td><td>Identification of the responsible party for placing the device into the market.</td></tr>
<tr><td>Address of the Manufacturer</td><td>Required for devices with the CE mark or UKCA mark.</td></tr>
<tr><td>Date of Manufacture</td><td>Not needed if embedded in the Serial Number.</td></tr>
<tr><td>Serial Number</td><td>Alpha-Numeric, 12-20 digits with the first 4 digits indicating:
Date of Manufacturing in Work Week and Year
WWYY1234567890123456.</td></tr>
<tr><td></td><td>[Europe] Compliance with EU WEEE directive 2010/19/EU.</td></tr>
<tr><td>CE</td><td>[Europe] Compliance with EU EMC directive 2014/30/EU and RoHS directive 2011/65/EU.</td></tr>
<tr><td></td><td>[Australia, New Zealand] Compliance with requirements of the relevant Australian ACMA Standards, under the Radiocommunications Act 1992 and the Telecommunications Act 1997.</td></tr>
<tr><td>VCCI</td><td>[Japan] Compliance with Japan VCCI requirements.</td></tr>
<tr><td>KC
XXXX-XX-XX</td><td>[Korea] Compliance with requirements of the Radio Research Laboratory Ministry of Information and Communication Republic of Korea.</td></tr>
<tr><td>CAN ICES-3(*)/NMB-3(*)</td><td>[Canada] Compliance with Canada standard ICES 003. Where * is either A or B.</td></tr>
<tr><td>FC</td><td>[USA] Optional. Compliance with United States Federal Communications Commission requirements.</td></tr>
</table></td></tr>
</table>

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><td>c UL us
XXXXX</td><td>[USA] Compliance with UL standards and Canadian Safety Standards.</td></tr>
<tr><td>XXXXX
RoHS</td><td>[Taiwan] Compliance with Taiwan EMC and RoHS.</td></tr>
<tr><td>20</td><td>[China] Compliance with Chinese environmental requirements. The number inside the circle is usually 10 or 20.</td></tr>
<tr><td></td><td>[Morocco] Compliance with Moroccan EMC standards.</td></tr>
<tr><td>UK
CA</td><td>[United Kingdom] Compliance with United Kingdom standards. Mandatory after 1/1/2022. Guidance to UKCA marking.</td></tr>
</table>
</td></tr>
<tr><td>LABL-16</td><td>If the surface of any component or casing will reach a temperature of 70°C (158°F) or greater the following warning logo shall be either printed on the label or placed separately on the device:
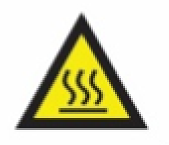
In addition, instructional safeguards shall be provided in accordance with IEC 62368-1 3rd Edition or later.</td></tr>
<tr><td>LABL-17</td><td>The format of the Serial Number shall be WWYYSerialNumber with no leading spaces (e.g., WWYY1234567890123456).</td></tr>
<tr><td>LABL-18</td><td>The Model Number shall have no leading spaces.</td></tr>
<tr><td>LABL-19</td><td>Barcodes shall be printed using Datamatrix ECC200 or Model 2 QR code only.</td></tr>
<tr><td>LABL-20</td><td>QR codes shall use a minimum of ECC Level M (15%).</td></tr>
<tr><td>LABL-21</td><td>The density of the barcode shall be 10 mil or larger.</td></tr>
<tr><td>LABL-22</td><td>The label shall only be printed on Polyester or Plastic labels using a Wax/Resin ribbon.</td></tr>
<tr><td>LABL-23</td><td>The PSID shall be printed on the label in its direct 32-character alphanumeric representation without any ancillary delimiting characters (e.g., underscore, dash, backslash, forward slash, etc.) and shall exactly match the readout of Barcode-2 (see the label example at the top of this section).</td></tr>
</table>

15 Compliance

15.1 ROHS Compliance

Requirement ID	Description
ROHS-1	The Supplier shall provide component-level reporting on the use of listed materials by concentration (ppm) for all homogenous materials.

15.2 ESD Compliance

Requirement ID	Description
ESD-1	Device manufacturer shall provide ESD immunity level (HBM- Human Body Model) measured in accordance with IEC-61000-4-2.

15.3 EMC Compliance

Requirement ID	Description
EMC-1	The device shall comply with the requirements of FCC / EN55032 Class B Radiated Emissions.

16 Shock and Vibration

Below are the shock and vibration specifications for storage devices:

Requirement ID	Description
SV-1	The non-operational shock requirement is 700G, half-sine, 0.5ms, total 6 shocks, along all three axes (+/-).
SV-2	The vibration requirement during operation is: 1.8G$_{rms}$, 5-500-5 Hz, Random Vibe, 20 min along all three axes.
SV-3	The vibration requirement during non-operation is: 3.13G$_{rms}$, 5-800-5 Hz, total 6 sweeps along all three axes, 20 minutes per sweep.
SV-4	Validation flow for Shock and Vibration: 1. UUT (Unit Under Test), test fixture should be visually inspected and ensured that everything is torqued or secured as needed. Pictures of test fixture with and w/o UUT should be provided. 2. Baseline performance of device should be gathered and used as a reference against post S&V data to ensure no performance impact incurred. 3. Once S&V testing is completed, repeat visual inspection to the UUT and test fixture to ensure no physical damage or performance impact has occurred to the UUT or test fixture. 4. Re-run stress test on the UUT in case of non-op test and provide data indicating no performance impact incurred to the unit.

17 Sustainability Requirements

The following are the sustainability requirements for devices.

Requirement ID	Description
SUS-1	A Life Cycle Assessment (LCA) aligned with ISO 14040 for the device shall be provided to the customer.
SUS-2	Recycled content information per material type shall be provided to the customer.

18 NVMe Linux CLI Plug-In Requirements

18.1 NVMe CLI Management Utility

The NVMeCLI utility (https://github.com/linux-nvme/nvme-cli/tree/master/plugins/ocp) shall be used as one of the management utilities for NVMe devices.

Requirement ID	Description
UTIL-1	The SSD supplier must test their SSDs with this utility and ensure compatibility. The following is the minimum list of commands that need to be tested with NVMeCLI: • Format. • Secure erase. • FW update. • Controller reset to load FW. • Health status. • Log page reads including vendor log pages. • SMART status. • List devices. • Get/set features. • Namespace management. • Identify controller and namespace. • Effects log page. • Latency Monitoring.

18.2 NVMe CLI Plugin Requirements

Requirement ID	Description
UTIL-PI-1 through UTIL-PI-7	Obsolete. See Section 18.1 NVMe CLI Management Utility.

18.2.1 NVMe CLI Plug-In Nomenclature/Functional Requirements

Requirement ID	Description
UTIL-NM-1 through UTIL-NM-10	Obsolete. See Section 18.1 NVMe CLI Management Utility.

18.2.2 NVMe CLI Plug-In FW Activation History Requirements

Requirement ID	Description
UTIL-FWHST-1 Through UTIL-FWHST-15	Obsolete. See Section 18.1 NVMe CLI Management Utility.

19 Revision History

Revision	Date	Release Notes
1.0	03/16/2020	Initial release based on feedback from the industry.
1.0a	06/26/2020	Errata.
2.0	05/18/2021	Additional major features and requirements including Latency Monitoring, Device Capabilities, Unsupported Requirements, Datacenter SSD Power States, Multiple Namespaces, Sanitize, NVMe-MI, Write Zeroes, Compare, Fused, Write Uncorrectable, Device Profiles, SPDM and additional security requirements, etc.
2.5	09/22/2023	Additional major features and requirements including Telemetry Data Area 1 & 2, Hardware Component log, FDP, DIX, security requirements, etc.

20 META Specific Items

The following items apply specifically to devices delivered to Facebook.

20.1 Configuration Specifics

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>FB-CONF-1</td><td>Obsolete.</td></tr>
<tr><td>FB-CONF-2</td><td>Obsolete.</td></tr>
<tr><td>FB-CONF-3</td><td>For all form factors, SMBus byte 91 bit 6, Firmware Update Enabled bit shall be set to 1b by default from the factory.</td></tr>
<tr><td>FB-CONF-4</td><td>Devices shall not support Error Injection (Feature Identifier C0h) Set Feature.</td></tr>
<tr><td>FB-CONF-5</td><td>Devices shall not support Error Injection (Feature Identifier C0h) Get Feature.</td></tr>
<tr><td>FB-CONF-6</td><td>Devices shall not support Error Recovery (Log Identifier C1h).</td></tr>
<tr><td>FB-CONF-7</td><td>Obsolete</td></tr>
<tr><td>FB-CONF-8</td><td>The default power state shall conform to the following table:
<table>
<tr><th>Form Factor</th><th>Capacity</th><th>Default Power State Upon Factory Exit</th></tr>
<tr><td rowspan="4">E1.S 25mm</td><td>1 TB</td><td>7 (10W)</td></tr>
<tr><td>2TB</td><td>6 (12W)</td></tr>
<tr><td>4TB</td><td>5 (14W)</td></tr>
<tr><td>8 TB</td><td>4 (16W)</td></tr>
<tr><td rowspan="2">E1.S 9.5mm</td><td>1 TB</td><td>8 (8.25W)</td></tr>
<tr><td>2 TB</td><td>7 (10W)</td></tr>
</table>
</td></tr>
<tr><td>FB-CONF-9</td><td>Devices shall be configured to Configuration Setting A as shown in Section 13 Device Profiles.</td></tr>
</table>

20.2 Performance Requirements

The following numbers are the Facebook performance targets for data storage SSD across all form factors. They are provided to serve as guidance for SSD Vendors. Performance scripts can be found on GitHub at *https://github.com/facebookincubator/FioSynth.* The targets are broken down into the following segments:

<table>
<tr><th>Requirement ID</th><th>Description</th></tr>
<tr><td>META-PERF-1</td><td>FB-FIO Synth Flash Targets (for all capacities).</td></tr>
<tr><td>META-PERF-2</td><td>fb-FIOSynthFlash TRIM Rate targets.</td></tr>
<tr><td>META-PERF-3</td><td>IO.go benchmark target.</td></tr>
<tr><td>META-PERF-4</td><td>Fileappend benchmark target.</td></tr>
<tr><td>META-PERF-5</td><td>Sequential write bandwidth.</td></tr>
<tr><td>META-PERF-6</td><td>Cache Bench target.</td></tr>
<tr><td>META-PERF-7</td><td>All targets shall be achieved by using “kyber” as the I/O scheduler.</td></tr>
<tr><td>META-PERF-8</td><td>The power measurement methodology is described in PCM-1, PCM-2, and PCM-3.
<table>
<tr><th>Form Factor</th><th>Capacity</th><th>Max Average Power Consumption to Achieve All Performance Targets</th></tr>
<tr><td>M.2</td><td>2TB and smaller</td><td>8.5W</td></tr>
<tr><td>E1.S</td><td>1TB</td><td>8.5W</td></tr>
<tr><td></td><td>2TB</td><td>10W</td></tr>
<tr><td></td><td>4TB</td><td>12W</td></tr>
<tr><td></td><td>8TB</td><td>14W</td></tr>
</table>
</td></tr>
<tr><td>META-PERF-9</td><td>BootBench Target</td></tr>
<tr><td>META-PERF-10</td><td>All tests shall be executed on FBK 5.19.0 FBK8 or later versions with support for io_uring engine.</td></tr>
<tr><td>META-PERF-11</td><td>Additional performance data shall be provided to Meta based on the following configuration in Figure 1.</td></tr>
</table>

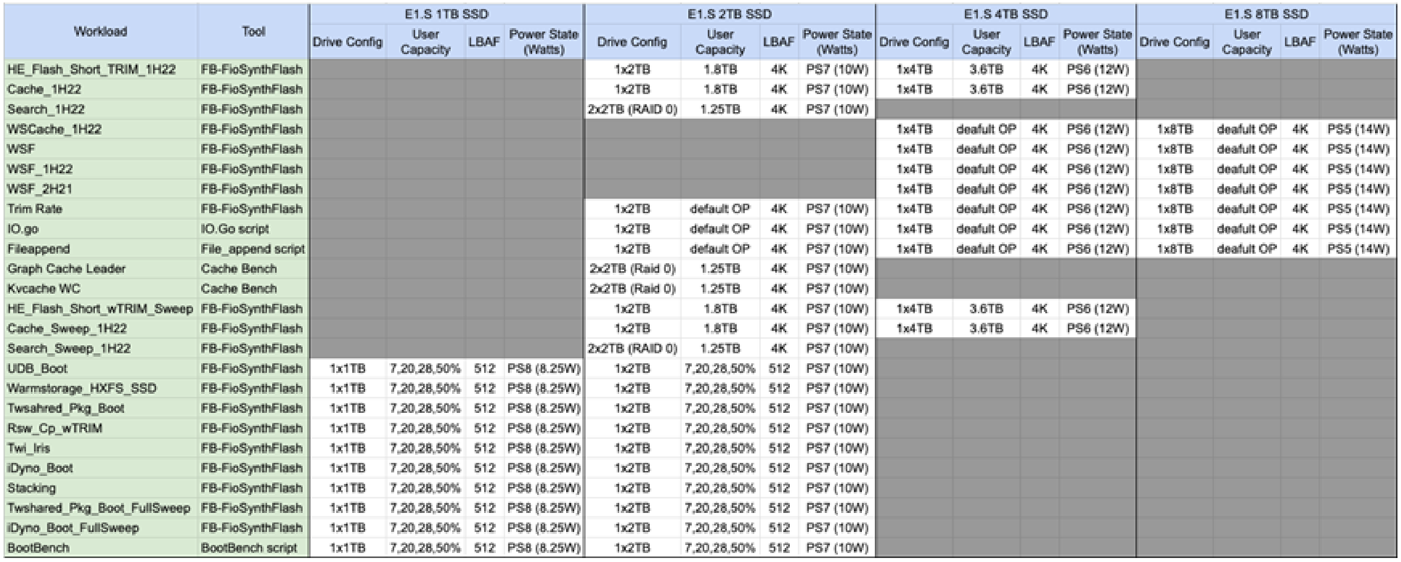

Workload	Tool	E1.S 1TB SSD				E1.S 2TB SSD				E1.S 4TB SSD				E1.S 8TB SSD			
		Drive Config	User Capacity	LBAF	Power State (Watts)	Drive Config	User Capacity	LBAF	Power State (Watts)	Drive Config	User Capacity	LBAF	Power State (Watts)	Drive Config	User Capacity	LBAF	Power State (Watts)
HE_Flash_Short_TRIM_1H22	FB-FioSynthFlash					1x2TB	1.8TB	4K	PS7 (10W)	1x4TB	3.6TB	4K	PS6 (12W)				
Cache_1H22	FB-FioSynthFlash					1x2TB	1.8TB	4K	PS7 (10W)	1x4TB	3.6TB	4K	PS6 (12W)				
Search_1H22	FB-FioSynthFlash					2x2TB (RAID 0)	1.25TB	4K	PS7 (10W)								
WSCache_1H22	FB-FioSynthFlash									1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
WSF	FB-FioSynthFlash									1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
WSF_1H22	FB-FioSynthFlash									1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
WSF_2H21	FB-FioSynthFlash									1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
Trim Rate	FB-FioSynthFlash					1x2TB	default OP	4K	PS7 (10W)	1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
IO.go	IO.Go script					1x2TB	default OP	4K	PS7 (10W)	1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
Fileappend	File_append script					1x2TB	default OP	4K	PS7 (10W)	1x4TB	deafult OP	4K	PS6 (12W)	1x8TB	deafult OP	4K	PS5 (14W)
Graph Cache Leader	Cache Bench					2x2TB (Raid 0)	1.25TB	4K	PS7 (10W)								
Kvcache WC	Cache Bench					2x2TB (Raid 0)	1.25TB	4K	PS7 (10W)								
HE_Flash_Short_wTRIM_Sweep	FB-FioSynthFlash					1x2TB	1.8TB	4K	PS7 (10W)	1x4TB	3.6TB	4K	PS6 (12W)				
Cache_Sweep_1H22	FB-FioSynthFlash					1x2TB	1.8TB	4K	PS7 (10W)	1x4TB	3.6TB	4K	PS6 (12W)				
Search_Sweep_1H22	FB-FioSynthFlash					2x2TB (RAID 0)	1.25TB	4K	PS7 (10W)								
UDB_Boot	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Warmstorage_HXFS_SSD	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Twsahred_Pkg_Boot	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Rsw_Cp_wTRIM	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Twi_Iris	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
iDyno_Boot	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Stacking	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
Twshared_Pkg_Boot_FullSweep	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
iDyno_Boot_FullSweep	FB-FioSynthFlash	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								
BootBench	BootBench script	1x1TB	7,20,28,50%	512	PS8 (8.25W)	1x2TB	7,20,28,50%	512	PS7 (10W)								

Figure 1 Performance Test Drive Configuration

20.3 Performance Targets for FB-FIO Synth Flash - HE_Flash_Short_wTRIM_1H22 (for all capacities)

Command: fb-FioSynthFlash -w HE_Flash_Short_wTRIM_1H22 -d ALL - f HE_Flash_Short_wTRIM_1H22_results

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
4K_L2R6DWPD_wTRIM	68 MB/s	72 MB/s	117 MB/s	1,500 us	2,500 us	5,000 us	10,000 us	8,000 us	15,000 us	20,000 us
4K_L2R9DWPD_wTRIM	68 MB/s	93 MB/s	156 MB/s	1,500 us	2,500 us	8,000 us	10,000 us	10,000 us	15,000 us	20,000 us
MyRocks_Heavy_wTRIM	210 MB/s	101 MB/s	22 MB/s	1,500 us	2,500 us	8,500 us	15,000 us	8,000 us	15,000 us	20,000 us
Fleaf	320 MB/s	87 MB/s	89 MB/s	1,500 us	2,500 us	10,000 us	15,000 us	12,000 us	15,000 us	25,000 us

20.4 Performance Targets for FB-FIO Synth Flash – Cache_1H22 (for all capacities)

Command: fb-FioSynthFlash -w Cache_1H22 -d ALL - f Cache_1H22

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
B_Cache	164 MB/s	96 MB/s	0 MB/s	1,000us	2,000 us	15,000 us	20,000 us	15,000 us	20,000 us	30,000 us

20.5 Performance Targets for FB-FIO Synth Flash – Search_1H22

Command: fb-FioSynthFlash -w Search_1H22 -d ALL - f Search_1H22

Workload	Read MB/s per Node	Write MB/s per Node	TRIM MB/s per Node	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
SearchLM_ wTRIM	2,550 MB/s	12 MB/s	130 MB/s	1,500 us	5,000 us	15,000 us	20,000 us	15,000 us	20,000 us	25,000 us

20.6 Performance Targets for FB-FIO Synth Flash – WSF

Command: fb-FioSynthFlash -w WSF -d ALL - f WSF #(2TB and 4TB capacities)

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
wsf-tl-heavy	237.6 MB/s	146.2 MB/s	87.7 MB/s	1,000 us	4,000 us	9,000 us	25,000 us	42,000 us	45,000 us	50,000 us

Command: fb-FioSynthFlash -w WSF_1H22 -d ALL -f WSF_1H22 #(8TB capacities)

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
wsf-tl-normal	78 MB/s	34 MB/s	29 MB/s	1,000 us	5,000 us	10,000 us	25,000 us	10,000 us	15,000 us	25,000 us

Command: fb-FioSynthFlash -w WSF_2H21 -d ALL -f WSF_2H21 #(8TB capacities)

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
wsf-tl-normal	112.4 MB/s	68 MB/s	41.4 MB/s	1,000 us	4,000 us	8,000 us	20,000 us	20,000 us	30,000 us	50,000 us

20.7 Performance Targets for FB-FIO Synth Flash – WACache_1H22 (8TB capacity)

Command: fb-FioSynthFlash -w WSCACHE_1H22 -d ALL -f WSCACHE_1H22 #(8TB capacities)

Workload	Read MB/s per TB	Write MB/s per TB	TRIM BW per TB	P99 Read Latency	P99.99 Read Latency	P99.9999 Read Latency	Max Read Latency	P99.99 Write Latency	P99.9999 Write Latency	Max Write Latency
WSCache_ 1H22	215 MB/s	90 MB/s	N/A	10,000 us	15,000 us	20,000 us	30,000 us	50,000 us	85,000 us	100,000 us

20.7.1 Performance Targets for Boot-Drive QOS workloads

Workload	Workload	IO Type	IOPS	B/W (MB/s)	P99 (ms)	P99.99 (ms)	Max Latency (ms)
UDB_Boot	Read	N/A	N/A	5	20	500	
	Write	N/A	N/A	100	200	400	
Warmstorage_HXFS_SSD		Read	N/A	N/A	5	20	500
		Write	N/A	N/A	40	100	400
Twsahred_Pkg_Boot		Read	210	51	60	90	100
		Write	600	62	30	60	250
Rsw_Cp_wTRIM		Read	119	27	15	30	80
		Write	700	80	30	60	120
		Trim	N/A	60	N/A	N/A	N/A
Twi_Iris		Read	420	45	10	50	150
		Write	65	6	40	200	400
		Trim	N/A	2.4	N/A	N/A	N/A
iDyno_Boot		Read	N/A	75	25	50	90
		Write	N/A	15	30	45	60
		Trim	N/A	0.1	N/A	N/A	N/A
Stacking		Read	N/A	120	10	50	120
		Write	N/A	64	0.6	60	350
		Trim	N/A	18	N/A	N/A	N/A

20.8 Trim Rate Targets

- This test measures raw trim performance with no background I/O
- 64M trim >= 50GiB/s & <= 10ms P99 trim latency
- 3GB trim >= 500GiB/s & <= 10ms P99 trim latency
- 4K time >= 10K trims/sec & <= 19ms P99 trim latency

20.9 IO.go Benchmark Targets

- This test measures how long the file system is blocked from writing/overwriting a file while a different file is deleted.
- Less than 4 file sizes total with latency outliers > 10ms
- No more than 2 latency outliers per file size
- No single latency outlier above 15ms
- The IO.go targets will be met with XFS, EXT-4 & BTRFS file systems.

20.10 Fileappend Benchmark Targets

- This test measures how long the file system is blocked from appending to a file while a different file is deleted.
- No measurable stalls reported by this tool.
- Max acceptable latency outlier is 10ms when deleting 1GiB or 2GiB file.
- Test should be executed for xfs, ext4 and btrfs file systems.

20.11 Sequential Write Bandwidth

- Full device (all available user capacity, all namespaces) must be written/filled in 180 minutes or less.
- Simple single-threaded sequential write FIO script to fill device.

20.12 CacheBench Targets

- A benchmarking tool that is a supplement for FB FIO Synth Flash tool on measuring performance for cache applications. This is different than the "B Cache" workload in FB FIO Synth Flash.
- Two workloads need to be tested:
  - Graph Cache Leader
  - Kvcache WC
- The final allocator and throughput stats from the benchmark will be used to see if the targets are met.
- Send SSD latency versus time file to Facebook using one of the following methods:
  - Send the raw results log file.
  - Run the "extract_latency.sh script and return the raw results log file, ".tsv" and ".png" files.
- Vendor NVMe CLI plug-in with "physical NAND bytes written" metric in the SMART / Health Information Extended (Log Identifier C0h) needs to be working to get the write amplification.

Workload	Get Rate	Set Rate	Read Latency (us)						Write Amp
			P50	P90	P99	P99.9	P99.99	Max	
Graph Cache Leader	87,000	16,000	100	300	800	N/A	3,000	12,000	1.3
Kvcache WC	3,200	1,500	300	700	2,000	9,000	14,000	15,000	1.4
Workload	**Get Rate**	**Set Rate**	**Write Latency (us)**						**Write Amp**
			P50	**P90**	**P99**		**P99.99**	**P100**	
Graph Cache Leader	87,000	16,000	30	50	100	N/A	700	8,000	1.3
Kvcahce WC	3,200	1,500	100	200	400	9,000	7,000	8,000	1.4

20.13 Max Latency Targets

This requirement aims to ensure the max latency for reads/writes are within certain bounds even as workload rate scales. We expect the workloads contained in this section to not exceed the max latency requirements defined here.

Workload	Max Read Latency (ms)	Max Write Latency (ms)
HE_Flash_Short_wTRIM_Sweep	65	85

Workload	Max Read Latency (ms)	Max Write Latency (ms)
Search_Sweep_1H22	65	85
Cache_Sweep_1H22	65	85
Twshared_Pkg_Boot_FullSweep	100	250
iDyno_Boot_FullSweep	90	60

20.14 BootBench Target

This test measures load on a Boot Device in a stacked environment.

- Download the tool from: github.com/liu-song-6/bootbench.
- Execute ./run.py . FIO 3.20 or newer must be installed as a prerequisite.
- Look at the final_result.txt file for output at the end of the run!

Workload	Read IOPS (sustained)	Overall Result
BootBench	At least 60K	Pass, without any errors.

21 Microsoft Specific Items

The following items apply specifically to Microsoft.

21.1 Configuration Specifics

Requirement ID	Description
MS-CONF-1	E1.S form factor devices and M.2 form factor devices shall be formatted to 512-byte sectors from the factory.
MS-CONF-2	E1.L form factor devices shall be formatted to 4096-byte sectors from the factory.
MS-CONF-3	Obsolete.
MS-CONF-4	Obsolete.
MS-CONF-5	Obsolete.
MS-CONF-6	For all devices, SMBus byte 91 bit 6, Firmware Update Enabled bit shall be set to 1b (Firmware Update is Enabled) by default from the factory.

22 Latency Monitoring Feature Set Theory of Operation

22.1 Overview

Latency outliers are very undesirable in the data center. The goal of this feature is to allow production monitoring of SSD QOS outliers and to debug outlier issues in production. This feature will allow suppliers and hyperscale companies to clearly understand if the outliers are caused by the SSD or other components in the Host system. This feature enables SSD suppliers to effectively debug latency issues efficiently. This feature will enable predicting when latency outliers are growing and likely to impact hyperscale customers.

This feature enables many use cases involving understanding and debugging latency issues at scale in a production environment.

22.2 Functional Operation

22.2.1 Latency Monitoring Feature Description

Below is the high-level theory of operation describing how the Latency Monitoring feature works. The latency of an individual command shall be measured from the time a controller fetches the command from the SQ to the time when controller writes the CQ entry for the SQ entry it fetched.

22.2.2 Bucket Description Overview

There are two types of buckets Active Buckets and Static Buckets. Active Buckets are buckets that are updated in real time. Static Buckets are buckets which are loaded with snapshots from the Active Buckets. This is a move from the Active Buckets to the Static Buckets, thus the old values in the Static Buckets are discarded. The Static Buckets allow hyperscale users to sample the Static Buckets over a fixed time period to gather statistics.

22.2.3 Active Bucket Description

The high-level concept is to create 4 real time buckets groups of active latency tracking command counters. Each bucket will count latency events which exceed a configured latency threshold. Below is a description of each bucket:

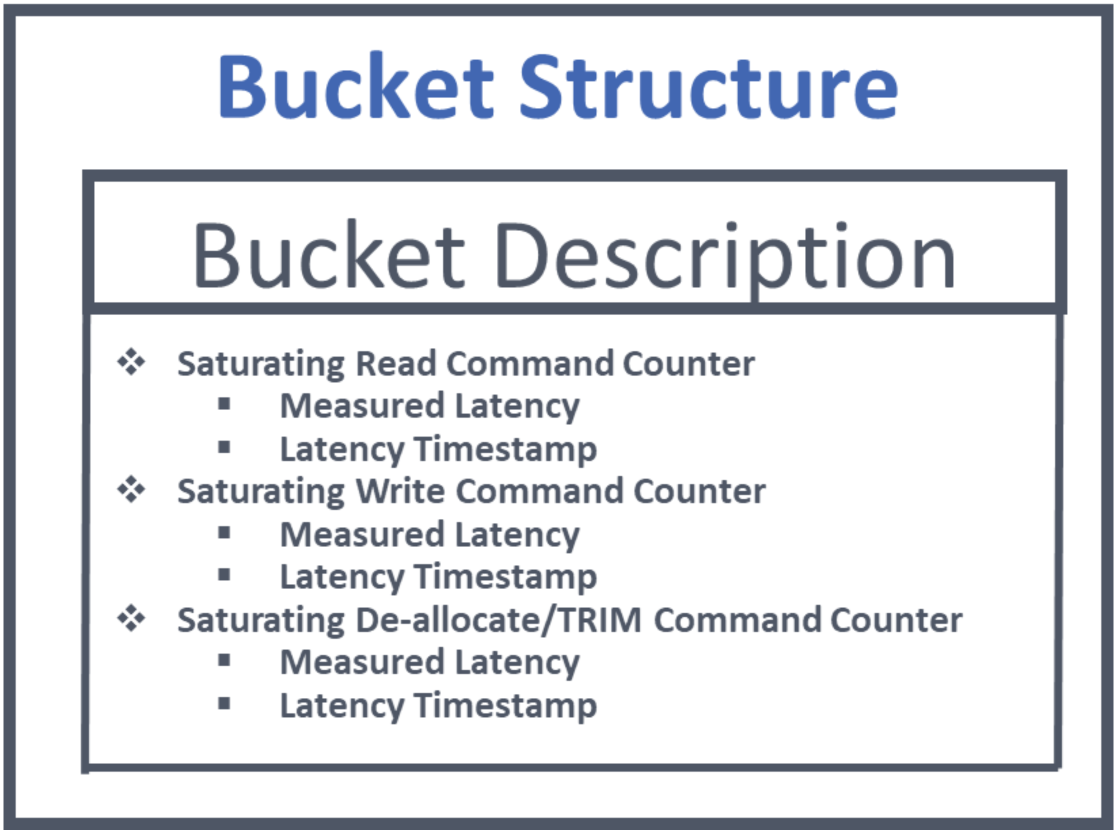

Bucket Structure

Bucket Description

- ❖ Saturating Read Command Counter
  - ▪ Measured Latency
  - ▪ Latency Timestamp
- ❖ Saturating Write Command Counter
  - ▪ Measured Latency
  - ▪ Latency Timestamp
- ❖ Saturating De-allocate/TRIM Command Counter
  - ▪ Measured Latency
  - ▪ Latency Timestamp

Each bucket contains the following:

- Saturating Read Command Counter with an associated Measured Latency and Latency Timestamp.
- Saturating Write Command Counter with an associated Measured Latency and Latency Timestamp.
- Saturating De-allocate/TRIM Command Counter with an associated Measured Latency and Latency Timestamp.

For clarity, the opcode to Counter mapping is below:

Bucket Counter	Opcode	Command
Read Command Counter	02h	Read
Write Command Counter	01h	Write
De-allocate/TRIM Command Counter	09h with Attribute – Deallocate (AD) = 1	Dataset Management

In addition to the command counters there is a Measured Latency data structure and a Latency Timestamp data structure associated with each command counter. The Measured Latency and Latency Timestamp have a direct relationship such that both are updated, or neither are updated. The Measured Latency and Latency Timestamp will be described later in this document.

22.2.3.1 Active Command Counter Behavior

The active command counters count Read commands, Write commands, and De-Allocate/TRIM commands which exceed a configured latency threshold. These active command counters count until the command counter saturates or the Active Bucket Timer expires. Below is the behavior for each of these events:

- Active Command Counter Saturation:
    - If the Command Counter saturates, the counter shall maintain the active value and not wrap.
- Active Bucket Timer Expiration:
    - If the Active Bucket Timer Expires, then the following occurs:
- The Active Bucket Command Counter values and associated information are moved into Static Bucket Command Counters, the Active Bucket Command Counters are then cleared to zero and re-start counting.
- The active command counters shall count regardless of how the Active Latency Minimum Window is configured.

22.2.3.2 Active Bucket Thresholds

There are multiple Active Command Counter Buckets. The configured latency thresholds determine which Bucket the command counter shall be incremented in. Below is a picture showing how there are multiple Buckets and thresholds for each Bucket.

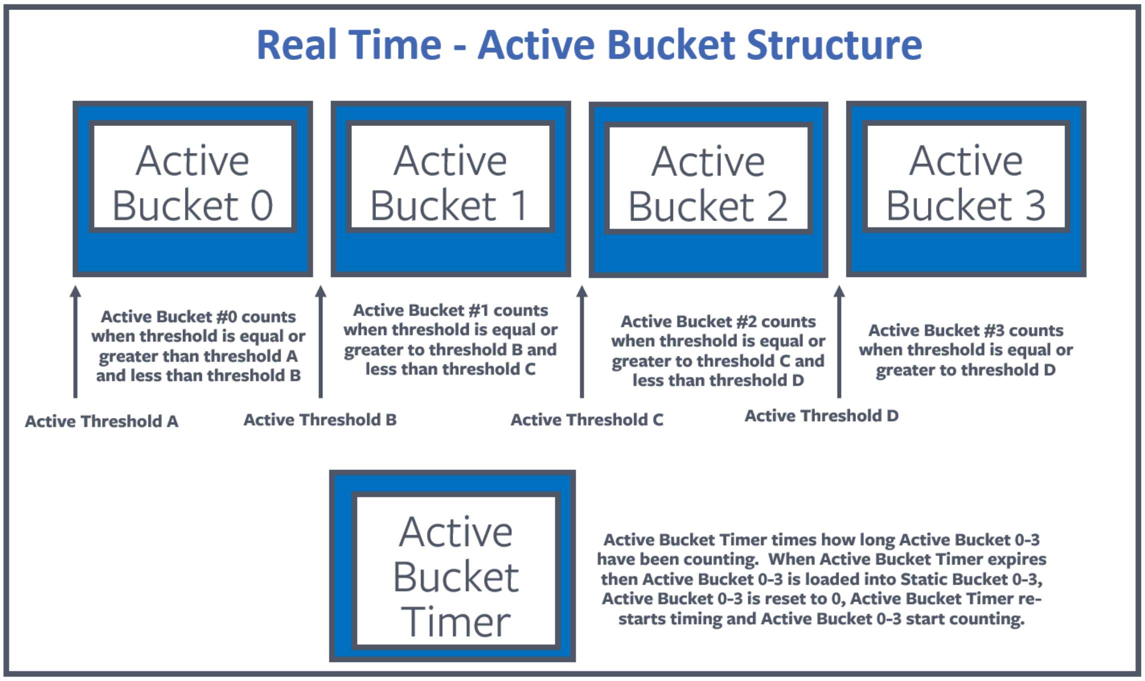
Real Time - Active Bucket Structure
Active Bucket 0
Active Bucket 1
Active Bucket 2
Active Bucket 3
Active Bucket #0 counts when threshold is equal or greater than threshold A and less than threshold B
Active Bucket #1 counts when threshold is equal or greater to threshold B and less than threshold C
Active Bucket #2 counts when threshold is equal or greater to threshold C and less than threshold D
Active Bucket #3 counts when threshold is equal or greater to threshold D
Active Threshold A
Active Threshold B
Active Threshold C
Active Threshold D
Active Bucket Timer
Active Bucket Timer times how long Active Bucket 0-3 have been counting. When Active Bucket Timer expires then Active Bucket 0-3 is loaded into Static Bucket 0-3, Active Bucket 0-3 is reset to 0, Active Bucket Timer re-starts timing and Active Bucket 0-3 start counting.

If the Read, Write, De-allocate/TRIM command completion time is below Threshold A then no counter is incremented. If the threshold is equal or greater than A and less than B, then the corresponding command counter in Active Bucket 0 increments. If it is equal to or greater than B and less than threshold C, then the corresponding command counter in Active Bucket 1 increments. If it is equal to or greater than C and less than threshold D, then the corresponding command counter in Active Bucket 2 increments. If it is equal to or greater than threshold D, then the corresponding command counter in Active Bucket 3 increments. By following this process all latencies greater than threshold A are counted for Read, Write and De-Allocate/TRIM commands. When configuring the Latency Monitor Feature the thresholds shall always be configured such that Active Threshold A < Active Threshold B < Active Threshold C < Active Threshold D.

22.2.3.3 Active Bucket Timer Behavior

The Active Bucket Timer times how long the Bucket Command Counters have been counting. When the Active Bucket Timer is equal to the Active Bucket Timer Threshold then the following operations shall occur:

1. The following data is moved:
    a. Active Bucket Counters 0 - 3 are moved to Static Bucket Counters 0 - 3.
    b. Active Latency Timestamps are moved to Static Latency Timestamps.
    c. Active Measured Latencies are moved to Static Measured Latencies.
    d. Active Latency Timestamp Units are moved to Static Latency Timestamp Units.
2. The Active Bucket items shall then be updated as follows:
    a. Active Bucket Counters 0 - 3 are cleared to zero.
    b. Active Latency Timestamps are set to invalid (FFFF_FFFF_FFFF_FFFFh).
    c. Active Measured Latencies are cleared to zero.
    d. Active Latency Timestamp Units are cleared to zero.
    e. The Active Latency Minimum Window, if it is running, may be reset or continue to count.
    f. Active Bucket Timer is cleared to zero and starts the process of counting over.

g. When looking at the data structures in the Latency Monitor (Log Identifier C3h) it should be noted that the data in item #1 and #2 above is 16-byte aligned, and the data can be moved simply by doing a data move of the entire data structure.

22.2.3.4 Active Latency Timestamp Format

The format of the Latency Timestamp follows the Timestamp data structure which is defined in the NVMe Specification. The Latency Timestamp allows an understanding of where a latency excursion occurred in terms of wall clock time. The Latency Timestamp time reported shall be based on CQ completion.

If the device receives a Timestamp Set Features (Feature Identifier 0Eh), then the device shall use this baseline wall clock time plus the time in milliseconds since the Timestamp was set to determine the Latency Timestamp of the latency event. If the device receives multiple Timestamp Set Feature (Feature Identifier 0Eh) the most recent Timestamp shall be used as the baseline.

If the device does not receive Set Features with a Timestamp, then the Latency Timestamp shall be generated based on Power on Hours.

If the device receives a Set Features with a Timestamp and then the device is powered off. When the device is powered on it shall use the most recent Timestamp it received even if this Timestamp was from before the device was powered off.

The Active Latency Timestamp Units shall be populated when the Latency Timestamp is updated to indicate if the Latency Timestamp used Timestamp with Power on Hours to generate the Latency Timestamp or if only Power on Hours were used to generate the Latency Timestamp.

22.2.3.5 Active Measured Latency and Active Latency Timestamp Updates

The Active Measured Latency is the latency measured from fetching the command in the SQ to updating the CQ. The Active Measured Latency data structure and the Active Latency Timestamp data structure shall be updated atomically. They shall not update independently. Each Command Counter (Read/Write/De-Allocate) has an Active Measured Latency and an Active Latency Timestamp structure associated with it. The Active Latency Configuration is used to configure this feature.

When the Latency Mode in the Active Latency Configuration is cleared to zero then the following behavior shall be followed:

- The Active Measured Latency and Active Latency Timestamp will be loaded the first time the command counter associated with it increments.
- The Active Measured Latency and Active Latency Timestamp will not be loaded again until the Active Measured Latency is reset.

If the Latency Mode is set to 0001h in the Active Latency Configuration, then every time the associated command counter is incremented the Active Measured Latency will report the largest measured latency based on the associated command counter. The Active Latency Timestamp will report the time when the largest latency occurred.

The Active Latency Minimum window also affects when the Active Measured Latency and the Active Latency Timestamp are updated. This is described in the section on Active Latency Minimum Window.

22.2.3.6 Active Latency Minimum Window

This affects both the Active Measured Latency and the Active Latency Timestamp. This defines the minimum time between updating the Active Measured Latency and the Active Latency Timestamp for a single Active Bucket/Counter combination. The feature is only used if the Active Latency Mode is set to 0001h.

If the Active Latency Minimum Window timer is running and the Measured Latency and Latency Timestamp have been updated, then the Latency Timestamp and Active Measured Latency will not be updated again until the Active Latency Minimum Window timer has expired. Below are some examples of this:

Example 1:
Assume:

- Active Latency Minimum Window of 5 seconds.
- Latency Mode is Configured for Largest Latency.
- Bucket 2 has a threshold range of 40ms to 400ms.

Time in seconds	Read Counter Bucket 2 Latency Event	Active Read Counter Bucket 2 Value	Actual Latency	Active Measured Latency	Latency Stamp	Comment
0	N	0	-	-	FFFF_FFFF_ FFFF_FFFFh	Actual Latency and Active Measured latency are invalid.
0.5	Y	1	50ms	50ms	0.5 Seconds	First Latency Event. This starts the Active Latency Minimum window. New latency events will not be recorded until the 5 second Active Minimum Window expires at 5.5 seconds.
5.25	Y	2	100ms	50ms	0.5 Seconds	Measured Latency and Latency Timestamp is not updated due to Minimum Window is not expired; however, the Active Read Counter is updated.
6	Y	3	75ms	75ms	6 Seconds	Minimum Window is expired and 75ms is greater than the previous number of 50ms, so the Active Measured Latency and Latency Timestamp are updated.

Example 2:
Assume:

- Active Latency Minimum Window of 5 seconds.
- Latency Mode is Configured for First Latency Event.
- Bucket 2 has a threshold range of 40ms to 400ms.

Time in seconds	Read Counter Bucket 2 Latency Event	Active Read Counter Bucket 2 Value	Actual Latency	Active Measured Latency	Latency Stamp	Comment
0	N	0	-	-	FFFF_FFFF_ FFFF_FFFFh	Actual Latency and Active Measured latency are invalid.
0.5	Y	1	150ms	150ms	0.5 Seconds	First Latency Event.
10	Y	2	200ms	150ms	0.5 Seconds	Since the device is in First Latency Event mode, no additional events are recorded.
15	Y	3	75ms	150ms	0.5 Seconds	Since the device is in First Latency Event mode, no additional events are recorded.

The Active Latency Minimum Window acts as a filter to ensure there are not a large number of events to update the Active Measured Latency and the Latency Timestamp. Thus, if the queue depth is 128 commands deep and there is a latency event, then there are not 128 updates to these data structures. Rather the first event is recorded, and the rest of the events are filtered out. It should be noted that the Active Latency Window is not enforced across power cycles. Thus, after a power cycle the Active Latency Window shall not start until a Bucket Counter is incremented.

22.2.3.7 Static Bucket Description

In addition to the real time buckets there are static buckets. When the Active Bucket Timer reaches its configured threshold (Active Bucket Timer Threshold) the active real time buckets shall be loaded into the static buckets and the active real time buckets shall be reset and start counting from 0. Below is a picture describing this:

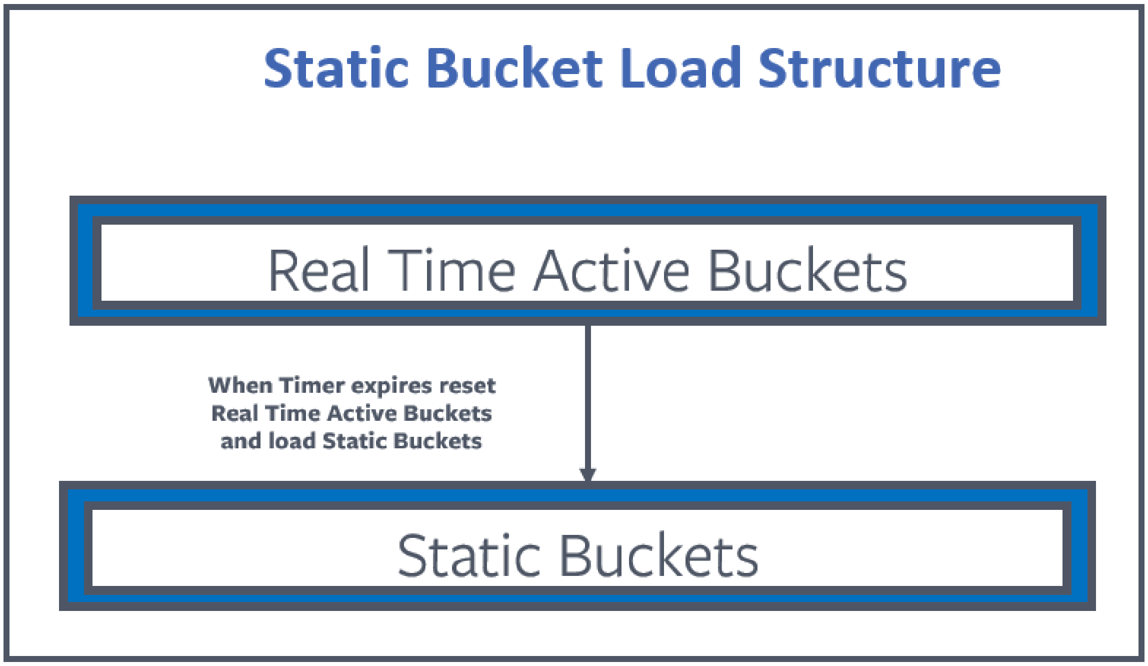
Static Bucket Load Structure
Real Time Active Buckets
When Timer expires reset Real Time Active Buckets and load Static Buckets
Static Buckets

22.3 Persistence Across Power Cycles

When either a safe or unsafe power transition happens the counters and associated Latency Monitoring Information shall be saved such that they can be restored on the next power up transition and continue from where the power cycle left off. There are additional details about this in the C.5 Latency Monitoring Feature - Challenging Event Handling section.

22.4 Debug Logs

The Latency Monitoring Feature can also enable debug logs to trigger. The Debug Log Trigger Enable configures which counters shall trigger a debug log the first time the Bucket/Counter combination is incremented. Only a single debug log shall be generated. Once a Latency Monitor debug log is generated, until the Latency Monitor debug log is discarded another Latency Monitor debug log cannot be generated. The Latency Monitor debug log shall be discarded using Set Features for the Latency Monitor or by reading the Latency Monitor Debug Log.

The Set Features for the Latency Monitor has two mechanisms for discarding the Debug Log. One method discards the debug log and resets the Latency Monitor feature to a new set of configured values based on the fields in Set Features. The other method discards the Debug Log and has no effect on any of the other features associated with the Latency Monitor Feature. Thus, the Latency Monitor Feature will keep running undisturbed when the Debug Log is discarded.

When the Latency Monitor debug log trigger event happens, the following data shall be captured: Debug Log Measured Latency, Debug Log Latency Timestamp, Debug Log Trigger Source, Debug Log Timestamp Units, Debug Log Pointer as well as internal information required to debug the issue to root cause.

22.5 Latency Monitoring Feature - Challenging Event Handling

22.5.1 Power Off/On When Latency Monitoring Feature is Enabled

When powering off, the Active Bucket Information (Counters, Measured Latency, Latency Timestamp, Active Bucket Timer) may be slightly off due to concerns with flushing data with unsafe power down. The Active Bucket Information shall maintain coherency compared to itself when flushing data with unsafe power down. When the device powers back on the Latency Monitoring Feature shall restore the Active/Static Bucket/Debug Information including loading the Active Bucket Counter. Once the restoration is complete then the device shall resume the Latency Monitor functionality. The device shall start capturing latency data within 2 minutes of power on. Thus, commands for the first 2 minutes may not be monitored. The Active Latency Window is not enforced across power cycles. Thus, after a power cycle the Active Latency Window shall not start until a Bucket Counter is incremented.

22.5.2 Power Off/On When Latency Monitoring Feature is Disabled

The Latency Monitoring Feature disabled/enabled state shall be persistent across power cycles. When the Latency Monitor Feature is enabled the associated Latency Monitor configuration information in the Latency Monitor Log shall persist across power cycles.

22.5.3 Moving Information from Active to Static Buckets

When moving data from the active buckets to static buckets it can be challenging to track all the information. The Active counter shall restart counting no longer than 3 seconds from moving data from Active to Static Counters.

22.5.4 Firmware Update

When activating new firmware, if the Latency Monitoring Feature is enabled, the firmware activation shall reset the Latency Monitoring Feature just as if a Set Features command to enable the feature was received. The Latency Monitoring Log shall start updating properly within 2 minutes of firmware activation completing. Thus, there are command latencies which could be missed after initially activating new firmware.

22.6 Configuring Latency Monitoring

The Set Features command is used to configure this feature. When the Set Features command is used to configure this feature all the data structures in both the Active and Static Buckets are reset. The debug log and related debug information is reset based on the Discard Debug Log Field.